ISSN 1936-5330

kcFED Research Working Papers

Reconciling VAR-based Forecasts with Survey Forecasts

Taeyoung Doh and A. Lee Smith

December 2018; Revised April 2020
RWP 18-13
https://dx.doi.org/10.18651/RWP2018-13

Federal Reserve Bank of Kansas City

Reconciling VAR-based Forecasts with Survey Forecasts*

Taeyoung Doh † A. Lee Smith‡

April 2020

Abstract

Expectations about future economic conditions play a central role in macroeconomic theory. Vector autoregression (VAR) models are often used to both generate forecasts to proxy these expectations and to model the dynamics of survey forecasts. However, jointly analyzing realized data and survey forecasts in a VAR remains a challenge. The primary issue that arises when embedding realized data alongside survey forecast in a VAR is the simultaneous existence of two different expectations of the same variable: the VAR-based forecast and the survey forecast. This paper proposes a Bayesian prior over the VAR parameters which allows the econometrician to impose the desired degree of consistency between these two forecasts at low computational costs. Our approach leverages the existence of multiple forecasts to aid in structural VAR identification and enhance VAR forecasts. We illustrate the usefulness of our approach in several applications including the identification of forward guidance shocks, the degree to which households' inflation expectations are anchored, and the evolution of inflation tail-risks during and after the Great Recession.

Keywords: Vector Autoregression (VAR), Survey Forecasts, Bayesian VAR, Monetary Policy, Forward Guidance, Inflation Expectations, Inflation Targeting, Inflation, Inflation Risk

JEL classification codes: C11, C32, E52, E31

*We are grateful for computational and data support from CADRE. We also thank participants at various conferences including: the 2018 FRB-Chicago & NBER Midyear Meeting of the DSGE Models and Applications Group and the 2019 System Econometrics Group Meetings. We also appreciate thoughtful comments from our formal discussants, Daniel Lewis and Paul Ho. Jason Bae and Logan Hotz provided outstanding research 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.

†Economic Research Department, Federal Reserve Bank of Kansas City, 1 Memorial Drive, Kansas City, MO 64198. Phone: (816) 881-2780, Email: Taeyoung.Doh@kc.frb.org.

‡Economic Research Department, Federal Reserve Bank of Kansas City, 1 Memorial Drive, Kansas City, MO 64198. Phone: (816) 881-2294, Email: andrew.smith@kc.frb.org.

1 Introduction

Modern macroeconomics emphasizes the role that expectations about the future play in shaping current economic conditions. For instance, the expected future path of interest rates — not just rates today — is most relevant for consumption and investment decisions. Similarly, expectations about future inflation are thought to be a primary driver of how much firms change their prices today. Within a model expectations are often modeled under assumptions that lead to a unique expectation, or forecast, for each variable. For instance, in a unique bounded rationale expectations equilibrium, there exists a single bounded forecast for each variable. This results in only one relevant measure of expected interest rates or expected inflation for agents' decision making. However, empirically examining the evolution of expectations outside the strict confines of a fully-specified model remains a challenge. In particular, there are several potential ways to obtain forecasts of a variable and therefore there exists potentially conflicting expectations for the same variable.

Two common approaches to measuring expectations are to either use forecasts based on a vector autoregression (VAR) as proxies for expectations (Campbell and Shiller, 1988; Keating, 1990) or to use available survey data as proxies for expectations of economic agents. Each approach is imperfect. For instance, using VAR forecasts as proxies for expectations assumes that the variables included in the VAR capture all relevant information available to agents, an assumption which may not hold when the VAR includes only a small number of variables. And survey forecasts lack in two dimensions. First, survey forecasts, in and of themselves, do not provide a data-generating process that can be used to analyze the formation and evolution of expectations. Second, available survey forecasts may not be the relevant expectations for economic decision making. The shortcomings of each approach can be mitigated by embedding survey forecasts into a VAR model. This dual approach greatly enriches the information set of parsimonious VAR models and enables formal analysis of the feedback effects between realized data and survey expectations without the need to assume that the available survey forecasts align perfectly with the expectations of economic agents. Therefore, a growing literature has included survey expectations directly in structural VAR models to either identify structural shocks or to analyze the dynamic interactions between expectations and macroeconomic aggregates in response to structural shocks.[1]

However, augmenting a VAR with survey data raises conceptual and computational challenges. Conceptually, without imposing consistency between the VAR forecast and the survey

[1]Examples include Leduc, Sill and Stark (2007); Clark and Davig (2011); Coibion and Gorodnichenko (2012); Sims (2012); Barsky and Sims (2012); Leduc and Sill (2013); Melosi (2017).

forecast, there exists two different forecasts for the same variable. Large inconsistencies between the VAR forecast and the survey forecast leaves the econometrician to wonder whether her model is misspecified or if the survey data is deficient.[2] Previous work has sought to resolve this tension by imposing restrictions on the VAR model's coefficients. Cogley (2005) provides one such example where expectations of the future short-term interest rate embedded in the current long-term interest rate are assumed to be consistent with the VAR forecasts of the short-term interest rate. However, this severely restricts the VAR dynamics by assuming that the expectations hypothesis holds. Therefore, imposing strict consistency between the two forecasts may require adherence to a particular theory such as the expectations hypothesis or full information rational expectations; these very theories may often be the test subject of the empirical investigation at hand, as in Coibion and Gorodnichenko (2012).[3] The approach in Kozicki and Tinsley (2012) imposes that expectations based on a time-varying intercept autoregressive model of inflation are, on-average, equal to survey-based measures of expected inflation. This relaxes the strict-consistency requirement in Cogley (2005) but still assumes that any difference between the VAR model and the survey forecasts is white noise. Neither approach is easily generalized and doing so can come at a high computational cost in larger VAR models with a rich lag structure.

In this paper, we propose a Bayesian approach to address the challenges posed by including realized data and survey forecasts of the same variable, or a closely related variable, in a VAR. In particular, we argue for constructing a non-degenerate prior for VAR coefficients that places greater mass on areas of the parameter space where the two forecasts align. We call this the *forecast consistent prior.* Our prior allows for dynamic correlation between the realized data and its' survey forecast, unlike the popular Minnesota prior, without dogmatically restricting the discrepancy between the two forecasts.[4] From an implementation standpoint, our method relies on a computationally efficient importance sampling technique which simply re-weights the posterior draws obtained without using the forecast consistent prior. Intuitively, under the forecast consistent prior, a draw which is closer to satisfying forecast consistency receives a greater weight. These weights are informed by a hyperparameter which governs the precision of our prior. By setting this parameter very small or very

[2]Many surveys are thought to be deficient in some manner. For example, Coibion et al. (2020) argue that firm surveys of inflation expectations ask only vague questions or are not nationally representative.

[3]There may also be measurement issues which call into question strict consistency. One example would be the differences between real-time and revised data. There may also be conceptual differences between the variable and the survey forecasts. For example, the Livingston Survey/Survey of Professional Forecasters provides forecasts of GNP growth prior to 1992 and GDP growth thereafter.

[4]The Minnesota prior ignores dynamic cross correlations by assuming each series follows a random walk.

large, our approach nests the cases of either no consistency or strict consistency between the VAR-based and survey forecasts.

A novel aspect of our prior is its applicability to structural VAR models to aid in the identification of structural shocks. For some intuition, note that structural impulse responses are simply conditional forecasts from the VAR. When the forecast consistent prior is placed on unconditional forecasts — in other words, forecasts that are not conditional on the realization of a structural shock and instead are based only on the observed data — then our prior shrinks the reduced-form VAR coefficients in the direction of satisfying forecast consistency. However, when the forecast consistent prior is placed on impulse response functions, or conditional forecasts, then our prior informs the structural VAR coefficients. One natural application which we explore in this paper is the addition of our prior to a structural VAR identified using sign-restrictions. Structural VAR (SVAR) models with only sign-restrictions on impulse responses lead to partial identification in which the model parameters are identified only up to a set (Moon and Schorfheide, 2012; Uhlig, 2017). Within this set, any two alternative SVAR models are equally probable. forecast consistency restrictions on impulse responses provide easily motivated probabilistic restrictions based on economic theory that can break this equality and distinguish two alternative SVAR models based on a variable's impulse response and the impulse response of its survey forecast.

We illustrate the usefulness of our approach in several settings to shed light on the effects of forward guidance shocks on output, the degree to which households' inflation expectations became better anchored after 2012, and the evolution of inflation tail-risks during the Great Recession and its aftermath. The forward guidance application is particularly illustrative. In this application we add Blue Chip forecasts of one-year ahead short-term interest rates to the Uhlig (2005) VAR model. We apply similar sign restrictions that Uhlig (2005) proposes and we find that, similar to Uhlig's finding for conventional monetary policy shocks, forward guidance shocks which reduce survey expectations of future short-term interest rates have an ambiguous effect on output. However, we also show that these sign-restrictions admit a wide range of time paths for the actual federal funds rate. Many of these time paths deviate significantly from the path of rates predicted by forecasters following the forward guidance shock. We then layer the forecast consistent prior on top of sign-restrictions to identify forward guidance shocks in which the VAR-implied path of the federal funds rate is tilted toward the path forecasters expect. We find that these shocks lead output to rise. These results suggest an important role for central bank communication to synchronize interest rate expectations in order to elicit the desired effects from forward guidance. Over our sample, we find that the Federal Reserve was typically able to achieve this synchronicity.

2 The Forecast Consistent Prior

In this section we introduce our VAR notation, next we use a simple example with a bivariate VAR to introduce the forecast consistent prior, and then we show how to apply our prior in a general VAR setting.

2.1 VAR Preliminaries

A reduced-form VAR(l) is given by:

$$Y_t = A_{(1)}Y_{t-1} + A_{(2)}Y_{t-2} + \ldots + A_{(l)}Y_{t-l} + u_t, \tag{1}$$

where Y_t is an $m \times 1$ vector of data at date $t = 1-l, \ldots, T$, $A_{(i)}$ are coefficient matrices of size $m \times m$, and u_t is the one-step ahead prediction error, or reduced-form residuals, which are assumed to be distributed normally with mean 0 and covariance matrix Σ. To simplify notation, we suppress means and other deterministic components of the VAR model. We assume that this reduced-from VAR is derived from an underlying structural VAR model:

$$BY_t = B_{(1)}Y_{t-1} + B_{(2)}Y_{t-2} + \ldots + B_{(l)}Y_{t-l} + \varepsilon_t, \tag{2}$$

in which B is an $m \times m$ non-singular coefficient matrix which governs the contemporaneous interactions between the variables Y_t, $B_{(i)}$ are coefficient matrices of size $m \times m$, and ε_t are the structural shocks which are independent of one another, mean zero, and have a standard deviation of one.

Combining equations (1) and (2) reveals that the reduced-form lag coefficient matrices $A_{(i)}$ are related to the structural coefficient matrices $B_{(i)}$ by the mapping $A_{(i)} = B^{-1}B_{(i)}$. And, similarly, the reduced form residuals u_t are related to the structural shocks by the mapping $u_t = B^{-1}\varepsilon_t$. Therefore, knowledge of B allows one to uncover the structural VAR model from the estimated reduced form VAR model. Following much of the structural VAR literature, we parameterize B as follows:

$$B^{-1} = CQ, \tag{3}$$

where C is the lower-triangular Cholesky factor of Σ and Q is a square $m \times m$ orthogonal rotation matrix such that $Q'Q = QQ' = I_m$. This parameterization is able to encompass multiple identification strategies including sign-restrictions which consider a set of alternative rotation matrices Q and recursive short-run restrictions which achieve point identification by assuming that $Q = I_m$.

We take a Bayesian approach to estimation and inference throughout the paper. In principle, our forecast consistent prior could be joined with any prior over the reduced-form VAR parameters. However, as we discuss later, popular priors such as the Minnesota prior may be undesirable in our setting where survey forecasts and realized data on the same variable are jointly included in Y_t. We describe the prior in terms of the stacked version of the reduced-form VAR in equation (1):

$$Y = XA + u, \tag{4}$$

where $X_t = [Y'_{t-1}, Y'_{t-2}, \ldots, Y'_{t-l}]'$, $Y = [Y_1, \ldots, Y_T]'$, $X = [X_1, \ldots, X_T]'$, $u = [u_1, \ldots, u_T]'$ and $A = [A_{(1)}, \ldots, A_{(l)}]'$. Within this notation, the OLS estimates of A, Σ, and C are given by:

$$\hat{A} = (X'X)^{-1}X'Y \tag{5}$$
$$\hat{\Sigma} = \frac{1}{T}(Y - X\hat{A})'(Y - X\hat{A}) \tag{6}$$
$$\hat{C} = Cholesky(\hat{\Sigma}). \tag{7}$$

We assume the following Normal and Inverse-Wishart conjugate priors for $\alpha = \text{vec}(A)$ and Σ, parameterized by $\nu_0, V_0, \alpha_0 = vec(A_0),$ and S_0:

$$\Sigma \sim IW(S_0, \nu_0),$$
$$\alpha|\Sigma \sim \mathcal{N}(\alpha_0, \Sigma \otimes V_0).$$

Given these priors, the posterior distributions of α and Σ become:

$$\Sigma|Y \sim IW(S_T, \nu_T), \tag{8}$$
$$\alpha|\Sigma, Y \sim \mathcal{N}(\alpha_T, \Sigma \otimes V_T), \tag{9}$$

where:

$$\nu_T = \nu_0 + T,$$
$$V_T = [V_0^{-1} + X'X]^{-1},$$
$$A_T = V_T[V_0^{-1}A_0 + X'X\hat{A}], \text{ and}$$
$$S_T = (Y - X\hat{A})'(Y - X\hat{A}) + \hat{A}'(X'X)\hat{A} + A_0'V_0^{-1}A_0 - A_T'(V_0^{-1} + X'X)A_T,$$

with $\alpha_T = vec(A_T)$. We assume weak or non-informative priors throughout the empirical applications in this paper by setting V_0 to the zero matrix and $\nu_0 = 0$. Using our earlier parameterization of $\Sigma = B^{-1}B^{-1\prime} = CQQ'C'$, these Normal and Inverse-Wishart priors can be described by $p(\alpha, C, Q) = p(\alpha, \Sigma) = p(\alpha|\Sigma)p(\Sigma)$.

2.2 Introducing the Forecast Consistent Prior

We now introduce the forecast consistent prior. For illustrative purposes, assume for a moment that we are interested in a bi-variate VAR(1) given by:

$$\begin{bmatrix} \pi_t \\ E_t^S(\pi_{t+1}) \end{bmatrix} = A \begin{bmatrix} \pi_{t-1} \\ E_{t-1}^s(\pi_t) \end{bmatrix} + u_t = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \begin{bmatrix} \pi_{t-1} \\ E_{t-1}^s(\pi_t) \end{bmatrix} + \begin{bmatrix} u_t^1 \\ u_t^2 \end{bmatrix}, \tag{10}$$

where π_t is a variable of interest such as inflation and $E_t^S(\pi_{t+1})$ is the one-period ahead forecast of π_t obtained from survey data. When left unconstrained, this model implies two diverging forecasts for the same variable: the VAR-based forecast and the survey forecast. From the first equation of the VAR, we can generate the time t VAR-based forecast for π_{t+1}:

$$E_t^{\text{VAR}}(\pi_{t+1}) = a_{11}\pi_t + a_{12}E_t^S(\pi_{t+1}). \tag{11}$$

The only way for the VAR and survey forecasts to always be consistent with one another, so that $E_t^{\text{VAR}}(\pi_{t+1}) = E_t^S(\pi_{t+1})$ for all possible realizations of $u_t = [u_t^1, u_t^2]'$, is to impose the following restrictions on the VAR coefficients:

$$g(A) = e_1'A - e_2' = [0, 0] \;\Leftrightarrow\; [a_{11}, a_{12} - 1] = [0, 0], \tag{12}$$

where e_i is an $m \times 1$ selection vector with a 1 in the i'th position and 0's elsewhere. However, imposing strict consistency at all times assumes that the two forecasts are conceptually identical, an assumption which may be difficult to defend. Consider for a moment the case of inflation and inflation expectations, an area of great interest to macroeconomists. It is common for researchers to splice together multiple surveys of professional forecasters and even interpolate semi-annual surveys into a quarterly or monthly series (Clark and Davig, 2011). Therefore, survey measures of inflation forecasts that researchers include in their empirical analysis almost certainly contain measurement error. Furthermore, the survey forecasts from professional forecasters are, at best, an imperfect proxy for the expectations of actual price setters (Coibion, Gorodnichenko and Kumar, 2018). While acknowledging these issues, completely ignoring the relevance of the survey forecasts for the VAR forecasts neglects the potentially useful information the econometrician possesses about the a-priori relationships between the variables in the VAR.

By adopting a Bayesian approach, we can vary the tightness of these cross-equation restrictions to allow for the possibility that survey forecasts and VAR-based forecasts are related to one another, albeit imperfectly. We can consider the following three cases which vary based on the degree of forecast consistency imposed:

1. **Strict Forecast Consistency**: $g(A) = [0, 0]$. Therefore, $a_{11} = 0$ and $a_{12} = 1$ is dogmatically imposed.

2. **Forecast Consistent Prior**: $g(A) \sim \mathcal{N}(0, (\lambda W)^{-1})$. The forecast consistent prior is centered on zero so that, on average, VAR-based forecast and survey forecasts are consistent with one another.[5] However, the two forecasts may deviate from each other from time to time. We specify the forecast consistent prior as a normal distribution which is the maximum entropy prior for $g(A)$ under the constraint that the first moment of $g(A)$ is zero and the second moment of $g(A)$ is $(\lambda W)^{-1}$.[6] Therefore, the size of the potential forecast deviations are governed by the weighting matrix W and the tuning parameter λ.

3. **No Forecast Consistency**: $g(A)$ is left unrestricted which is equivalent to a diffuse prior over $g(A)$.

Strict forecast consistency and no forecast consistency can be regarded as limiting cases when λ approaches ∞ and 0, respectively. In terms of A, the prior density function of $g(A)$ can be treated as the likelihood function for A using observations satisfying these restrictions.[7]

We can calculate the posterior density of α under the forecast consistent prior at low computational cost by importance sampling. First, define the posterior density of α obtained without imposing forecasting consistency by $p(\alpha|Y, C)$ and define the forecast consistent prior

[5] Note, this assumption is not necessary. For instance, if one wanted to model the fact that, over say a training sample, survey forecasts are somehow biased then a non-zero mean could be included in the forecast consistent prior.

[6] Our choice of the maximum entropy prior is motivated by the fact that, in the information-theoretic sense, it minimizes the amount of prior information on higher-order moments of $g(A)$. See, for example, Robert (2007); Cover and Thomas (2012).

[7] The maximum entropy prior for $g(A)$ is similar to the limited information likelihood for A in Kim (2002) if first and second moment restrictions for $g(A)$ are interpreted as nonlinear moment restrictions for A. We replace sample moment conditions for artificial observations satisfying weak consistency by population moment conditions, as in Del Negro and Schorfheide (2004).

density kernel for α by $h(\alpha) = p(g(vec(A)))$. Second, obtain the forecast consistent posterior density of α:

$$p(\alpha|Y, C, g(A)) = \frac{h(\alpha)p(\alpha|Y, C)}{\int h(\alpha)p(\alpha|Y, C)d\alpha}. \tag{13}$$

Using the above definition, we can simulate posterior draws of α from $p(\alpha|Y, C, g(A))$ simply by re-weighting the posterior draws from $p(\alpha|Y, C)$. Define the following importance weight for α^d that is randomly drawn from $p(\alpha|Y, C)$, by:

$$w(\alpha^d) = \frac{h(\alpha^d)}{\sum_{j=1}^{M} h(\alpha^j)}, \; \alpha^j \sim p(\alpha|Y, C). \tag{14}$$

We then re-sample these draws according to the weights $[w(\alpha^1), \ldots, w(\alpha^M)]$ to simulate the posterior density $p(\alpha|Y, C, g(A))$.

The forecast consistent prior can be used on top of any fully specified prior distribution for the VAR coefficients. One widely used prior for α is the Minnesota prior in which each variable is centered around a univariate random-walk process so that, in the above example, α_0 is set to $[1, 0, 0, 1]'$. Therefore, the Minnesota prior ignores the cross-equation linkages between variables by setting prior means of off-diagonal terms of A to 0. Although this prior might be a good benchmark when there is no a-priori obvious dynamic correlation between variables, the Minnesota prior is not ideal when there is a clear dynamic linkage between VAR variables, as is likely to be the case when survey forecasts are included alongside realized data.[8]

2.3 The Forecast Consistent Prior in a Structural VAR Model

The previous section motivates our forecast consistent prior in the setting of a reduced-form VAR. In this environment, the forecast consistent prior offers a theoretically grounded approach to parameter shrinkage. However, the full conceptual appeal of this prior is best illustrated in the context of a structural VAR model. In a structural VAR, the forecast consistent prior informs the structural VAR parameters for which, even with infinite observations, the data is uninformative. Therefore, the benefits of the forecast consistent prior

[8]In the above bi-variate VAR(1) example, imposing the weak consistency prior on top of the Minnesota prior will induce multiple modes for the prior distribution of $[a_{11}, a_{12}]$, with the one from the Minnesota prior $(1, 0)$ and the other from the forecast consistency prior $(0, 1)$. The practice would be equivalent to imposing a mixture of two prior distributions, in which mixing weights are determined by the relative magnitude of the prior tightness of each prior.

can extend beyond shrinkage when the econometrician is attempting to identify structural innovations.

As discussed above, the mapping between the reduced form and structural VAR models is defined by the coefficient matrix B which we parameterize by $B^{-1} = CQ$ where C is the lower-triangular Cholesky factor of Σ and Q is a square $m \times m$ orthogonal rotation matrix. For illustrative purposes, consider for a moment the VAR(1), which can generalized to a VAR(l) by writing the VAR in companion form:

$$Y_t = AY_{t-1} + u_t = AY_{t-1} + CQ\varepsilon_t, \tag{15}$$

where ε_t are the structural innovations of interest. Suppose that we are interested in the effects of one particular structural shock which, without loss of generality, we label ε_t^1. Then the initial or impact-effect of a 1 unit realization of ε_t^1 is given by the first column of CQ which is simply the 1-step ahead forecast of Y_t conditional on $\varepsilon_t^1 = 1$ in period t and 0 in all other periods:

$$E_{t-1}^{VAR}(Y_t|\varepsilon_t^1 = 1) = CQe_1. \tag{16}$$

The $h-$step ahead impulse response to this structural shock can be written as:

$$IRF(A, C, Q, h|\varepsilon_t^1 = 1) = E_{t-1}^{VAR}(Y_{t+h}|\varepsilon_t^1 = 1) = AE_{t-1}^{VAR}(Y_{t+h-1}|\varepsilon_t^1 = 1) = A^hCQe_1, \tag{17}$$

for $h = 0, 1, \ldots, H$. The elements A and C of the h-step ahead impulse response can be informed by the observed data contained in Y. However, the matrix Q must be identified from prior economic reasoning. In other words, the likelihood function of the VAR is invariant to alternative choices of Q.

When the VAR contains both realized and survey data, forecast consistency provides theoretically grounded restrictions on the matrix Q which, together with C, provides the mapping between reduced form VAR residuals and structural shocks. For example, if we suppose once again that $y_t = [\pi_t, E_t^S(\pi_{t+1})]'$, where $E_t^S(\pi_{t+1})$ is the 1-step ahead forecast of π_t obtained from survey data, then survey consistency suggests the restrictions:

$$e_1'IRF(A, C, Q, h|\varepsilon_t^1 = 1) = e_2'IRF(A, C, Q, h-1|\varepsilon_t^1 = 1)$$

$$\Leftrightarrow$$

$$[e_1'A - e_2'I_m]A^hCQe_1 = 0, \tag{18}$$

should hold for $h = 2, \ldots, H$. Equation (18) reveals that, in general, the forecast consistent prior will shape the posterior distribution of A, C, and Q.[9] More formally, the forecast consistent prior can be expressed as $g(A, C, Q|H) \sim \mathcal{N}(0_{H-1}, (\lambda W)^{-1})$ where W is a $H - 1 \times H - 1$ weighting matrix and λ calibrates the overall tightness the econometrician wishes to impose on these forecast consistent restrictions.

Structural VAR identification in the presence of survey data can benefit from the use of the forecast consistent prior by down-weighting posterior draws of A, C and Q which imply impulse responses with divergent VAR-based and survey forecasts. While this is true regardless of the identification strategy pursued, a growing structural VAR literature aims to identify structural shocks of interest by sign restrictions which restrict the shapes of the impulse responses to identify parameters in Q. VAR models identified by sign restrictions typically find a large set of Q matrices that are compatible with these restrictions (Moon and Schorfheide, 2012; Uhlig, 2005, 2017). In this setting, forecast consistent priors may be especially of interest. In particular, the sign-restricted VAR literature has been criticized on the grounds that sign-restrictions identify an entire set of structural models, some of which may imply impulse responses which are inconsistent with plausible empirical specifications or equilibrium models (Arias, Caldara and Rubio-Ramirez, 2018; Wolf, 2018). Our approach provides theoretically motivated restrictions to narrow the set of potential structural VAR models based on the dynamic linkages between two different forecasts of the same variable.

We briefly illustrate the potential for our conditional forecast consistent prior to sharpen the identification of structural VAR parameters in the context of sign restrictions. First, we define the identified set of orthonormal matrices Q that satisfy the sign restrictions, denoted by $\mathcal{Q}$, and defined as follows:

$$\mathcal{Q} = \{Q|B^{-1} = CQ, e_i' \cdot IRF(\hat{A}, \hat{C}, Q, h|\varepsilon_{1,t} = 1)_r \geq (\leq) 0 \, , \forall r = 1, \cdots, R\}, \tag{19}$$

where, we assume for a moment that, A and C are fixed at their OLS estimates, $e_i' \cdot IRF(\hat{A}, \hat{C}, Q, h|\varepsilon_{1,t} = 1) \geq (\leq) 0$ denotes the r-th restriction on the impulse response of VAR variables and R is the total number of restrictions on impulse-responses, and where e_i is a selection vector with a 1 in the i'th location and zeros elsewhere. Since the data do not provide additional information to distinguish different Q matrices in the set of $\mathcal{Q}$, the conditional posterior of Q is same as the conditional prior, which is uniform over $\mathcal{Q}$. In other words, without further restrictions, each matrix from $\mathcal{Q}$ is equally probable, meaning $p(Q_i) = p(Q_j)$ for any Q_i and Q_j belonging to $\mathcal{Q}$. Our forecast consistent prior breaks this

[9] At the extreme of strict consistency, $e_1 A = e_2 I_m$, implying no restrictions on CQ.

symmetry by penalizing structural coefficients that generate the greater divergence between the VAR-based forecast and the survey forecast. Hence, our new prior for Q induces a new prior for the structural VAR matrix B given C, which is $\widehat{p}(B) \propto w(Q)p(B) \propto w(Q)$, where $w(Q)$ can be constructed based on $g(Q|A, C, H)$.

To make matters concrete, consider the simple VAR(1) model:

$$\begin{bmatrix} \pi_t \\ E_t^S(\pi_{t+1}) \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \pi_{t-1} \\ E_{t-1}^s(\pi_t) \end{bmatrix} + \begin{bmatrix} 1 & 0 \\ 0.5 & 1 \end{bmatrix} \begin{bmatrix} q_{11} & q_{12} \\ q_{21} & q_{22} \end{bmatrix}, \tag{20}$$

where $E_t^S(\pi_{t+1})$ is the 1-step ahead forecast of π_t obtained from survey data. We aim to identify a "news" shock for which survey forecasts increase. In addition to this sign normalization, we impose the sign restriction that the impact effect of this news shock on the realized data in π_t is positive. Identifying this shock requires requires identifying a column of Q. Without loss of generality, we label the news shock ε_t^1 and therefore aim to identify the elements of the first column of Q, $[q_{11}, q_{21}]'$. Without identifying restrictions, the identified set of q_{11} and q_{21} is the entire unit circle. We impose the following restrictions to further narrow the identified set of q_{11} and q_{21}.

1. **Normalization Restriction: The "news" shock increases survey forecasts in period t.**
$$R1(q_{11}, q_{21}) = \left\{ (q_{11}, q_{21}) \middle| \frac{\partial E_t^S(\pi_{t+1})}{\partial \varepsilon_t^1} = 0.5q_{11} + q_{21} > 0 \right\}.$$

2. **Sign Restriction: Realized data increases alongside survey forecasts in period t.**
$$R2(q_{11}, q_{21}) = \left\{ (q_{11}, q_{21}) \middle| \frac{\partial E_t^{VAR}(\pi_{t+1})}{\partial \varepsilon_t^1} = q_{11} > 0 \right\}.$$

3. **Forecast Consistency Restriction: The response of realized data in period $t + 1$ is consistent with the increase in survey forecasts in period t.**

$$R3(q_{11}, q_{21}) = \left\{ (q_{11}, q_{21}) \middle| \frac{\partial E_t^{VAR}(\pi_{t+1})}{\partial \varepsilon_t^1} - \partial \frac{E_t^S(\pi_{t+1})}{\partial \varepsilon_t^1} = q_{21} - 0.5q_{11} = 0 \right\}.$$

Under the forecast consistent prior, the forecast consistency restrictions is loosely imposed by assuming that $q_{21} - 0.5q_{11}$ follows a Normal distribution centered around 0. We denote this prior by $g(Q|A, C, H) = q_{21} - 0.5q_{11} \sim \mathcal{N}(0, \lambda^{-1})$.

We can graphically illustrate how the forecast consistent prior shapes the posterior set of the structural VAR parameters. Panel (a) of Figure 1 highlights the identified set of

(q_{11}, q_{21}) when we impose only the normalization and sign restrictions ($R1$ and $R2$). The blue line on the unit circle outlines the identified set under these two restrictions. We cannot discriminate different locations in the blue line without additional identifying restrictions. When we augment $R1$ and $R2$ with $g(Q|A, C, H)$, the forecast consistent prior, we can discriminate different points on the blue line in panel (a) in a probabilistic way. Panel (b) of Figure 1 illustrates the forecast consistent prior's impact on the posterior distribution of (q_{11}, q_{21}) by shading with darker colors the region of the parameter space that has a higher probability mass under our forecast consistent prior. The forecast consistent prior adds curvature to the posterior distribution of the space of structural parameters which enables the econometrician to discriminate between alternative structural VAR models that are all equally likely according to the sign-restrictions.[10] In our forward guidance application, which we turn to next, we illustrate in a much richer structural VAR model — identified by sign restrictions — how the forecast consistent prior can enhance structural identification.

3 Identifying Forward Guidance Shocks

The textbook view in macroeconomics is that what matters for consumption and investment decisions is the entire path of expected future interest rates, not just interest rates today (Woodford, 2003). This notion underpins the theoretical result that, through its ability to communicate a path of future short-term interest rates, a central bank retains powerful ammunition to combat economic downturns in the face of constraints on the current policy rate (Eggertsson and Woodford, 2003). While the Federal Open Market Committee (FOMC) has been offering rate guidance since at least 1994 when they began to issue post-meeting statements, the FOMC made greater use of forward guidance in recent decades. FOMC rate guidance extended over longer horizons and became especially explicit after the target federal funds rate reached its lower bound in late 2008. Policymakers now expect forward guidance to remain a key element of the Federal Reserve's policy toolkit to combat future recessions (Yellen, 2016).

[10]Notice the fact that our approach adds curvature to the posterior distribution of Q addresses another shortcoming of the sign-restriction literature which is the tendency for researchers to report point-wise median impulse responses across the draws from $\mathcal{Q}$ as a measure of the central tendency of the posterior distribution of impulse responses. As Kilian and Lütkepohl (2017) point out, the standard approach of reporting the point-wise median impulse response not correspond to any particular structural VAR model. In our approach, the posterior mode of $g(Q|A, C, H)$ is one particular structural model and provides an alternative measure of central tendency around which to perform inference.

Despite the FOMC exercising forward guidance more often and with greater specificity in recent decades, questions continue to linger about the effects of forward guidance on real economic activity. Identification of forward guidance shocks through high-frequency event studies suggests that more accommodative rate guidance eases financial conditions, which most economists and policymakers would associate with high levels of economic activity (Gurkaynak, Sack and Swanson, 2005). However, researchers who have analyzed the effects these high-frequency forward guidance shocks have on survey forecasts often find that when the Federal Reserve has communicated a more accomodative future path of interest rates, forecasters revise down their expectations for growth and employment. Campbell et al. (2012) and Nakamura and Steinsson (2018) attribute these perverse estimates to a "Fed information effect" by which central bank communication of future accommodation is tied to a bleaker economic outlook. The disconnect between high-frequency changes in financial conditions and lower frequency movements in survey forecasts provides little clarity to policymakers regarding the efficacy of forward guidance in promoting macroeconomic objectives.

Given the range of macroeconomic outcomes associated with forward guidance, our central question surrounds the response of output following a forward guidance shock. Forward guidance shocks are inherently shocks to expectations for future short-term interest rates; therefore we augment a standard monetary VAR model with survey forecasts of future interest rates. While it is crucial to include expectations of future policy rates in the VAR to identify shocks to interest-rate expectations, as Ramey (2016) stresses, embedding survey forecasts of future interest rates alongside the federal funds rate generates two simultaneous — and likely contradictory — forecasts for future interest rates: the VAR forecast and the survey forecast.

Our novel contribution to the forward guidance literature is to leverage the existence of these two interest rate forecasts together with our forecast consistent prior to aid in the identification of forward guidance shocks. To illustrate the role that our forecast consistent prior plays in shaping the estimated response of output following a forward guidance shock, we first estimate the structural VAR model using a pure sign restrictions approach following Uhlig (2005). We find that forward guidance shocks identified solely with sign restrictions have an ambiguous effect on output. And, in several instances, identified shocks which satisfy the sign restrictions implies that output persistently contracts following what is thought to be an expansionary forward guidance shock. On the other hand, when we add to the sign restrictions the forecast consistent prior, we find expansionary effects on output for large enough values of λ which governs the tightness of the forecast consistent prior.

Our key result is that forward guidance which synchronizes various interest rate expectations on a lower path of rates leads to higher levels of economic activity. Alternatively, forward rate guidance which shifts some interest rate expectations but not others may have an ambiguous or even perverse effect on output. Therefore, it remains an empirical question as to whether or not past FOMC communication has achieved the necessary synchronicity across interest rate expectations in order to elicit the desired effects from forward guidance. We present evidence that for a-priori plausible degrees of dispersion across VAR and survey forecasts of interest rates, output has modestly increased following FOMC communication which signaled a lower path of future interest rates. In particular, we select the value of λ to maximize the correlation between the forward guidance shock series from our SVAR and high-frequency measures of forward guidance shocks. The chosen value of λ implies that output expands following FOMC communication which is sufficiently clear to uniformly lower financial market, survey, and VAR-implied expectations of future interest rates.

3.1 Data and VAR Model

To estimate the effects of forward guidance shocks on output, we closely follow the model specified in Uhlig (2005). In particular, we use monthly GDP as produced by Macroeconomic Advisers, the consumer price index, an index of commodity prices, the effective federal funds rate, non-borrowed reserves, total reserves, and 1 year ahead Blue Chip consensus economic forecasts for the 3-month Treasury bill rate. We take 100 times the natural log of all variables except for the federal funds rate and Blue Chip forecast for the 3-month Treasury bill rate. Unlike Uhlig (2005), we use the CPI directly as opposed to using the CPI and PPI to interpolate the quarterly GDP deflater to a monthly frequency. Also key is the addition of survey forecasts of short-term interest rates. The inclusion of forecasted interest rates is central to our analysis of forward guidance shocks.

We model these time series as a VAR(3) where the number of lags is selected based on standard information criteria over the sample January 1994 to December 2007.[11] We begin our estimation in 1994 due to the chronology of communication from the Federal Open Market Committee (FOMC). Our interest lies in understanding the output effects of forward guidance, which has often been primarily issued through press statements at the conclusion of FOMC meetings. Prior to 1994, the FOMC did not issue press statements, limiting the starting point of our sample. We end our sample in 2007 due to the dynamics of non borrowed reserves. In particular, beginning in January of 2008, this series begins

[11]We use pre-sample data beginning in October of 1993 to account for the 3 lags included in the VAR.

to take on negative values which precludes using the natural log of this series after 2007. However, in a robustness check, we "correct" the non-borrowed reserves series to offset the fact some of the borrowing through the Fed's liquidity facilities appear to have been counted in borrowed reserves but not total reserves and extend our estimation sample through the zero lower bound period.

3.2 Sign Restrictions & The Forecast Consistency Prior

We follow Uhlig (2005) closely and identify a forward guidance shock by restricting the sign of the impulse response of commodity prices, the price level, and forecasts of future interest rates for the first 6 months after a forward guidance shock. The notion that forward guidance shocks cause nominal interest rates and prices to move in opposite directions is shared by standard sticky-price models (Eggertsson and Woodford, 2003), models which attribute a large role to a "Fed information effect" (Nakamura and Steinsson, 2018), and models which dampen the output effects of forward guidance through a discounted Euler equation (McKay, Nakamura and Steinsson, 2016). Therefore, we consider structural VAR models that result in inflation and future interest rates persistently moving in opposing directions to distinguish the monetary policy innovation from other demand shocks. However, depending on the reaction of the central bank, supply shocks, including productivity and commodity price shocks, could also cause inflation and nominal interest rates to move in opposite directions. However, according to the estimates in Barsky, Basu and Lee (2015) using U.S. data, both TFP shocks and TFP news shocks have historically caused inflation and short-term interest rates to comove. Kilian and Lewis (2011) similarly find that over our sample oil price shocks continue to influence inflation but there is no evidence of systematic monetary policy responses to oil price shocks. This suggests that we are on fairly solid ground by assuming that monetary policy shocks are the only innovations that result in inflation and nominal interest rates persistently moving in opposite directions.[12]

In addition to sign restrictions, we also employ the forecast consistent prior to identify forward guidance shocks. In particular, the VAR contains both the effective federal funds rate and survey forecasts of the 3-month Treasury Bill rate, 1 year ahead. While these financial instruments are conceptually distinct, forward guidance shocks may facilitate arbitrage activities to reduce whatever gap exists between the two interest rates. Moreover,

[12] We make no explicit orthogonalization to distinguish conventional monetary policy shocks from forward guidance shocks. However, we note that by restricting our focus to policy shocks which result in persistent movements of expected future interest rates as implied by 1 year ahead survey forecasts, we rule out many short-lived monetary policy shocks.

forward guidance which is uniformly understood should reduce dispersion across measures of interest-rate expectations. This motivates us to impose some degree of forecast consistency when identifying forward guidance shocks. Therefore, we apply our forecast consistent prior between the survey data on the 3-month Treasury bill rate and the VAR forecast of the federal funds rate. Specifically, we form a prior as follows:

$$\begin{aligned} g(Q|A,C,h) &= e'_{BC}IRF(A,C,Q,h-12|\varepsilon_1^{fg}=1) \\ &\qquad - \frac{1}{3}\sum_{j=0}^{2}\left[e'_{FF}IRF(A,C,Q,h-1+j|\varepsilon_1^{fg}=1)\right] \\ &= \frac{\partial}{\partial\varepsilon_1^{fg}}\mathbb{E}_{h-12}^{BC}R_h^{T-bill} - \frac{1}{3}\sum_{j=0}^{2}\left[\frac{\partial}{\partial\varepsilon_1^{fg}}FF_{h-1+j}\right], \end{aligned} \tag{21}$$

where e_{BC} and e_{FF} are selection vectors which respectively extract the responses of BlueChip forecasts and the federal funds rate and $\frac{\partial y_h}{\partial\varepsilon_1^{fg}}$ denotes the period h impulse response of variable y to a forward guidance shock that occurred in period 1. $\mathbb{E}_t^{BC}R_{t+12}^{T-bill}$ denotes the Blue Chip consensus economic forecast for the 3-month Treasury bill 4-quarters from now. In the monthly Blue Chip survey, forecasters report what they expect the 3-month T-bill to average over in the three months ending 4 quarters ahead. So, in December, forecasters report what they expect the yield on the 3-month Treasury bill to average in the three months of October, November, and December of the following year. [13] To the extent that the yield on the 3-month Treasury bill rate closely tracks the federal funds rate, as has been the case historically, then this forecast should be linked with what the average federal funds rate over the three months ending in December. The average federal funds rate over the three months ending 4-quarters from now, conditional on a forward guidance shock in period 1, is given by $\frac{1}{3}\sum_{j=0}^{2}\left[\frac{\partial}{\partial\varepsilon_1^{fg}}FF_{h-1+j}\right]$.[14]

In the context of forward guidance, higher values of λ imply greater synchronocity of interest rate forecasts since it will tend to drive the posterior distribution of the VAR and survey forecasts of interest rates together. However, alternative values of λ result in structural VAR models that fit the data equally well. To be clear, this issue arises in all structural VAR models that are not over-identified. For example, in the case of recursive short-run restrictions, there is no statistical criterion on which to select one candidate Cholesky ordering

[13] Depending on whether the month is at the end or beginning of the quarter, the horizon of the forecast varies between 12 and 14 months. However, we must fix the horizon when implementing forecast consistency. This is yet another reason why forecast consistency should be expected to only weakly hold.

[14] We can stack the forecast consistency restriction in equation (21) for $h = 13, 14, \ldots, H$ to form the forecast consistent prior $g(Q|A,C,H) \sim \mathcal{N}(0_{H-12}, I_{H-12}\lambda^{-1})$ where λ tunes the precision over the forecast consistent prior. In practice, we set H = 60 = 48+12 to encompass the 48 period impulse responses we plot.

over another. Analogously in this application, there is no obvious statistical reason to prefer small versus large values of λ. Instead, we first highlight how alternative values of λ affect the structural response of output. We then argue that given the effect that different values of λ have on the impulse response of output, the effectiveness of FOMC communication in synchronizing interest rate expectations is a salient factor shaping the estimated effects of forward guidance shocks. We then use external information from high-frequency event studies of FOMC announcements to tune λ and find that, tuning the forecast consistent prior based on this external information, output has modestly increased following FOMC communication which signals a lower path of future interest rates.

3.3 Implementation of Sign Restrictions & The Forecast Consistency Prior

We specify our reduced form VAR as:

$$y_t = A_0 + A_1 y_{t-1} + A_2 y_{t-2} + A_3 y_{t-3} + u_t, u_t \sim \mathcal{N}(0, \Sigma),\ A = [A_0, A_1, A_2, A_3]' \tag{22}$$

We specify a non-informative prior, as described in Section 2.1, by setting V_0 to the zero matrix and $\nu_0 = 0$. Therefore, the posterior distributions for $\alpha = vec(A)$ and Σ, the VAR lag coefficients and the covariance matrix, are centered at the OLS estimates. Using our earlier parameterization of Σ from Section 2.1, $\Sigma = CQQ'C'$, where where C is the lower-triangular Cholesky factor of Σ and Q is a square $m \times m$ orthogonal rotation matrix. The mapping between the reduced form VAR residuals u_t and structural VAR shocks ε_t, is governed by the linear mapping $u_t = CQ\varepsilon_t$.

To focus on the issue of identification, rather than inference, we initially keep $A = \hat{A}$ and $\Sigma = (Y - X\hat{A})'(Y - X\hat{A})/T$ and consider only random draws of Q. However, in a robustness check, we sequentially draw from the posterior of A and Σ as well as Q. We draw random orthogonal rotation matrices Q by drawing a 6×6 random square matrix denoted by χ, with each element of χ independently drawn from a standard normal distribution, and then we take the QR decomposition of χ using MATLAB's `[Q,R]=qr(`χ`)` function. Each draw of Q represents a candidate structural VAR model. Structural VAR models are kept if they satisfy the sign restrictions and are otherwise discarded.

After accumulating M structural VAR models, or equivalently, M orthogonal rotation matrices Q, which satisfy the sign restrictions, then we calculate the posterior weight under the forecast consistent prior according to:

$$w(Q(m)) = \frac{exp(-0.5\lambda g(Q(m)|A,C,H)'g(Q(m)|A,C,H))}{\sum_{k=1}^{M}\exp(-0.5\lambda g(Q(k)|A,C,H)'g(Q(k)|A,C,H))}. \tag{23}$$

These weights form the importance sampling weights we use to simulate the posterior set of impulse responses. In practice, we set $M = 5000$ and find that for our calibrated values of λ, the effective sample size – defined as $\hat{M} = (\sum_{k=1}^{M} w(B(k))^2)^{-1}$ – suggests the posterior weight is distributed across thousands of draws. Finally, we scale the size of each draw to have the same initial effect on forecasted interest rates before calculating the importance sampling weights. This prevents larger shocks from being penalized simply due to the size of the forward guidance shock. This closely follows the approach taken in Uhlig (2005, pg. 413).

3.4 Impulse Response Functions

Figure 2 shows the median and 68 percent error bands of impulse response functions across the draws that satisfy the forward guidance sign restrictions. In this figure we set $\lambda = 0$. Per the imposed restrictions, forecasted interest rates decline and commodity prices along with the overall price level rise for the first 6 months. In the months that follow the restricted periods, survey forecasts of interest rates remain low and commodity prices remain elevated. The price level continues to gradually climb throughout the impulse response horizon. Recall that our sign restrictions leave non-borrowed reserves, total reserves, and the path of the actual federal funds rate unconstrained. However, 12 months after the forward guidance shock, measures of bank reserves increase which precipitates a decline in the actual federal funds rate. After reaching a trough around the 1 year horizon, the federal funds rate begins rising back to its pre-shock level and thereafter continues to rise, resulting in a persistent interest rate overshoot.

How does the VAR-implied path of the federal funds rate compare to forecasters' expectations of short-term interest rates? The red-dashed line in the bottom-left panel shows the VAR-implied forecast for the 1 year ahead short-term interest rate based on the impulse response of the federal funds rate. One year after the forward guidance shock, actual short-term interest rates are lowered by an amount similar to what forecasters anticipated. However, in subsequent months, the path of short-term interest rates exceeds the path anticipated by forecasters and remains above the survey forecast for interest rate for several years.

The identified response of output using only sign restrictions suggests that forward guidance may not be effective in stimulating economic activity as output initially declines after the forward guidance shock and then gradually increases towards its pre-shock path. However, in this sign-restriction identified VAR, the VAR-implied path of interest rates is more restrictive than the path of interest rates anticipated by professional forecasters. This raises the question of whether the apparent inability of forward guidance to stimulate output stems from the disconnect between interest rate expectations.

To understand the effects of FOMC forward guidance which synchronizes interest rate forecasts, we now impose the conditional forecast consistent prior in conjunction with sign restrictions. That is to say that we now consider $\lambda > 0$ whereas in Figure 2 we set $\lambda = 0$. Of course, this requires choosing a value for λ. Before picking a particular value for λ, we show how the response of output is influenced by the choice of λ. To first gain some intuition for how alternative values of λ will influence the posterior distribution of the impulse responses, it is instructive to examine two particular candidate SVAR models. Among the 5000 draws that satisfy the sign restriction, we isolate two particular draws: the draw that comes the closest to satisfying the forecast consistency restriction in equation (21) and the draw that is the furthest from satisfying the forecast consistency restriction in equation (21). In other words, these are the two draws that register the highest and lowest values in our forecast consistent prior density. We respectively refer to these as the "best" and "worst" draws from a forecast consistency standpoint.[15]

Figure 3 plots the best and worst draws. The green-dashed-dotted line represents the best draw and the red-dashed line represents the worst draw. To visually understand what is behind these rankings, the top-right panel of Figure 3 shows the cumulative forecast deviation. By construction, the period 48 cumulative deviation is closer to zero for the best draw than for the worst draw. In other words, the best draw results in a path of the federal funds rate which more closely mirrors the path of rates that forecasters expected. For the worst draw, the path of the federal funds rate meaningfully diverges from the path expected by forecasters in a direction that suggests policy ends up being much more restrictive than survey forecasters anticipated. Although the response of output played no role in our selection, the best draw implies a persistent expansion in output while the worst draw suggests that output persistently declines following a reduction in the path of short-term interest rates. Given the distinct paths of interest rates that underlie each of these candidate SVAR models, the resulting differences in their real effects may not be surprising. However, the

[15]The "best" draw corresponds closely with the posterior mode over the set $\mathcal{Q}$.

diverging output responses in Figure 3 offer some evidence that forward guidance has the potential to be effective when it is able to synchronize interest-rate expectations.

The impulse response of output is heavily influenced by the choice of λ. A tighter forecast consistent prior, achieved by selecting higher values of λ, will place greater weight on draws like the "best" one under which output expands and lower weight on draws like the "worst" one under which output contracts. Figure 4 specifically focuses on the effect λ has on the response of output 18 months following the forward guidance shock. For smaller values of λ output fails to meaningfully rise — and may actually decline — following what ought to be an expansionary forward guidance shock. However, as λ increases so too does the response of output at 18 months. Beyond a certain threshold, the output response at 18 months ceases to increase with further increases in λ. Interpreting λ as governing the degree to which forward guidance announcements synchronize interest rate expectations, then Figure 4 provides empirical evidence that forecast agreement regarding FOMC communications is a salient factor shaping the estimated effects of forward guidance shocks. In particular, this figure makes clear that some minimum level of synchronicity in rate expectations is needed for forward guidance to be effective.

Which of these responses best characterizes the U.S. experience with forward guidance? To answer this, we must select a particular value of λ. As previously discussed, the data used to estimate the reduced-form VAR parameters cannot distinguish between alternative values of λ. Therefore, we bring external information to bear to calibrate λ. In particular, we select λ to maximize the correlation between our structural VAR forward guidance shocks and high-frequency financial market measures of forward guidance shocks as constructed by Gurkaynak, Sack and Swanson (2005). The path factor constructed by these authors is a measure of forward guidance shocks which are orthogonal to unexpected changes in the target federal funds rate. λ is selected from a grid of values across which we compare the correlation between the structural shocks from the SVAR with the updated path factor series from Bundick and Smith (2019). This approach to calibrating λ is most natural in our setting as it effectively extends the notion of synchronizing interest-rate expectations across VAR-based forecasts and survey forecasts to include some degree of synchronicity with the forecasts implied by financial markets.

By incorporating high-frequency financial market measures of monetary policy shocks into the identification of forward guidance shocks in our monthly SVAR model, we can relate our approach to the burgeoning external instruments approach to VAR identification of monetary policy shocks, as employed by Gertler and Karadi (2015). However, our approach is less exposed to the critiques in Ramey (2016) since we directly incorporate forward-looking survey measures of interest rate expectations into our VAR model which lessens concerns around VAR invertibility. Our approach is also related to the narrative sign restrictions approach proposed by Antolín-Díaz and Rubio-Ramírez (2018) in that it allows the econometric to incorporate narrative information to select among alternative SVAR models. However, as we show with an alternative calibration in Section 3.5, narrative or external information is not strictly required to implement our approach.

Figure 5 shows the median and 68 percent error bands of impulse response functions across the draws that satisfy the sign restrictions when reweighted with $\lambda > 0$. The top-left panel of Figure 5 shows that output rises in a gradual but persistent manner following an expansionary forward guidance shock for the value of λ which maximizes the correlation between the SVAR shocks and the high-frequency forward guidance shock series. The bottom left panel of Figure 5 illustrates that this setting of the forecast consistent prior reduces the deviation between the VAR-impled path of the federal funds rate and the path of rates anticipated by professional forecasters following a forward guidance shock. More precisely, the cumulative deviation between the VAR-based forecast and the Blue Chip forecast is reduced by roughly 25 basis points under under our calibration of λ. Therefore, when $\lambda > 0$, more of the anticipated interest-rate accommodation remains in place which appears to support the persistent expansion of output.

3.5 Impulse Response Functions for Alternative Specifications

The finding that forward guidance shocks which better synchronize interest rate expectations, as calibrated through larger values of λ, lead to more expansionary output effects is a common finding across several alternative VAR specifications. We consider three variants of our baseline VAR model: one which extends the estimation sample to include the recent zero lower bound period, one which calibrates λ without the use of high-frequency measures, and one which sequentially draws from the posterior of A and Σ as well as Q. We discuss each in turn.[16]

The recent zero lower bound period from 2008-2015 included more extensive use of forward guidance with the FOMC's conventional policy tool constrained. However, our 1 year ahead survey forecasts became truncated around 2010, according to Swanson and Williams (2014), therefore we leave this full sample estimation as a robustness check. The first column of Figure 6 shows the impulse responses when we extend the estimation sample through

[16]We include the full impulse responses for these alternative specifications in the appendix.

December 2015. Over this extended sample, we continue to find expansionary effects for the value of λ which best correlates the SVAR forward guidance shocks with high-frequency forward guidance shocks and contractionary effects when $\lambda = 0$.

The use of high-frequency measures of forward guidance shocks provides a natural approach to calibrating λ in our setting since it effectively extends the notion of forecast synchronicity from survey and VAR-based interest-rate forecasts to include some degree of synchronicity with financial-market measures of expected future interest rates. However, this approach to tuning λ may be restrictive since our monthly VAR model can, in principle, capture forward guidance shocks outside of regularly scheduled FOMC meetings which comprise the event set of the high-frequency forward guidance surprises. Therefore, in the second column of Figure 6, we show impulse responses when we calibrate λ to match the discrepancy between the VAR and survey forecasts for future interest rates over the range of sign-restricted impulse responses. This discrepancy then serves as the basis for the prior variance we expect between the VAR and survey forecasts without referencing high-frequency shock measures. Under this alternative calibration of λ we continue to find that output expands following a downward revision to forecasts of interest rates.

Finally, we can extend our forecast consistent prior to shape the full set of VAR parameters, including the posterior distributions of A and Σ, not just the set of rotation matrices $\mathcal{Q}$. More concretely, we sequentially draw from the posterior distributions of A, C where $\Sigma = CC'$, and $\mathcal{Q}$. A given joint draw d of the triplet $(A, C, Q)^d$ is kept if the associated impulse responses satisfy the sign restrictions and otherwise discarded. Then we re-weight this posterior set according to the forecast consistent prior $g(A, C, Q|H)$, as extended from equation (21).[17] The third column of Figure 6 shows the impulse responses jointly drawn from the posterior of the full set of VAR parameters. The range of responses is understandably wider compared to the baseline VAR model. However, the ranking of the median responses remains with larger values of λ implying more expansionary effects from forward guidance.

4 The Anchoring of Consumer Inflation Expectations

The evolution of inflation and inflation expectations has long been a focal point for central bankers tasked with price stability. The advent of inflation targeting, and the macroeconomic benefits it is thought to confer, has only further increased the attention that policymakers

[17] In the terminology laid out in Uhlig (2005), this is similar to the "pure sign-restrictions approach" whereas our baseline implementation is more closely related with his "penalty-function approach."

give these variables. Of particular importance is the degree to which inflation expectations are "anchored." Indeed, much of the intellectual capital behind the Federal Reserve's decision to adopt a numerical inflation target in January 2012 was developed in research that showed foreign central banks with a formal inflation target had better anchored inflation expectations (Gürkaynak, Levin and Swanson, 2010).

Research regarding whether the adoption of a formal inflation target in the U.S. successfully anchored inflation expectations has found mixed results depending on whose inflation expectations are analyzed. Bundick and Smith (2018) find that financial market measures of inflation compensation became better anchored sometime between 2009 and 2012, which spans the period between the addition of longer-run inflation to the FOMC's quarterly Summary of Economic Projections and the adoption of a 2% longer-run inflation target. Doh and Oksol (2018) find similar evidence that professional forecasters' long-term inflation expectations showed signs of better anchoring beginning around 2010. Detmeister et al. (2015) also find that the FOMC's announcement of an explicit inflation objective led professional forecasters' long-run inflation expectations to coalesce around 2 percent. However, these authors as well as Binder (2017) find that the 2012 adoption of an inflation target had little influence on households' longer-term inflation expectations. Coibion and Gorodnichenko (2015*b*) similarly argue that households' near-term inflation expectations are not well anchored and instead they strongly vary with energy prices.

We leverage our forecast consistent prior to more formally assess the degree to which households' longer-term inflation expectations have become better anchored since 2012 using a structural VAR model. Intuitively, if household's longer-term inflation expectations are anchored, they should not respond to short-lived, transitory impulses to inflation. Therefore, in contrast to previous research, we study the degree to which unanticipated innovations to CPI inflation spillover to households' expectations for 5- to 10-year inflation as measured in the University of Michigan's Survey of Consumers.[18] The timing of CPI data releases relative to when the household survey data are collected naturally lends itself to a recursive VAR identification strategy. Therefore, this application illustrates our forecast consistent prior in the context of a point-identified VAR model. In this application we also demonstrate an alternative approach to setting the hyperparameter λ that governs the forecast consistent prior in a setting where we lack external or narrative information.

[18] Our analysis is related to Binder (2017) who studies the convergence of individual household expectations towards 2 percent after 2012. However, instead of assessing the level at which inflation expectations might be anchored, we aim to assess the degree of anchoring.

While we find significant discrepancies between household and VAR-based inflation forecast before 2012, our results suggest that after 2012 households' longer-term inflation expectations became both better aligned with VAR-based expectations and better anchored. For example, prior to 2012, increases in energy prices led household's to increase their expectations for inflation over the next 5 to 10 years. However, after 2012, households' long-term inflation expectations are insensitive to swings in energy prices. We find similar changes in the pass through from shocks to food and core inflation after 2012. In contrast to households' inflation expectations, long horizon inflation forecasts implied by the VAR model suggests little variation in inflation expectations in response to CPI innovations both before and after 2012. Therefore, this application highlights the flexibility of our approach which, in contrast to dogmatically imposing strict forecast consistency, allows for household survey expectations to be governed by something other than our VAR model.

4.1 Data and VAR Model

We analyze impulse responses from a four variable structural VAR model to detect a potential change in the degree to which household's longer-term inflation expectations are anchored since 2012. We include CPI energy inflation and CPI food inflation in addition to overall CPI inflation to account for the documented fact that some consumer prices, such as gasoline and grocery prices, are more important than others in shaping household inflation expectations (Coibion and Gorodnichenko, 2015*b*; Cavallo, Cruces and Perez-Truglia, 2017). We add to these CPI inflation measures the median 5- to 10-year inflation expectation from the University of Michigan's Survey of Consumers.

We model these series as a VAR(3) over two distinct samples: January 2004 - December 2011 and January 2012 - December 2019 on the basis of lag-selection criteria which recommend 1 or 2 lags depending on the sample and the criteria. The start date of the recent sample period, January 2012 - December 2019, is dictated by the Federal Reserve's January 2012 adoption of an inflation target. It is precisely this change in FOMC communication that we wish to analyze. The start date of the early sample period, January 2004 - December 2011, is selected to provide a symmetric sample size. However, we find similar results if we extend the start date of the early sample to January 1990.

4.2 Shock Identification

The timing of CPI data releases relative to when the household survey data are collected naturally lends itself to identification by zero short-run restrictions. In particular, we use a recursive identification strategy with the following ordering: CPI energy inflation, CPI food inflation, the median of households' expectations for inflation over the next 5 to 10 years, and CPI inflation. This ordering allows us to identify 3 inflation shocks: an energy inflation shock, a food inflation shock, and a core inflation shock. Our recursive identification strategy aims to identify these shocks based on the following sequence of surveys and data releases.

We assume that households are aware of energy and food inflation in the current month when they respond to the University of Michigan survey. After all, the typical household frequents gasoline stations, grocery stores, and restaurants at least once a week. Therefore, our shock identification allows for inflation in salient goods, namely gasoline and food, to influence their inflation expectations within the current month. This argues for ordering energy and food inflation ahead of households' expectations for inflation over the next 5 to 10 years in our recursively identified VAR model. We distinguish energy from food inflation shocks by ordering energy inflation ahead of food inflation. Our assumption is that an exogenous increase in energy prices can spillover to food prices within the month, perhaps through transportation and delivery costs. Conversely, we assume that an exogenous increase in food inflation can't spillover to energy inflation within the month.

Finally, we identify a core inflation shock on the basis of the timing of the consumer interviews conducted by the University of Michigan. Every month, the University of Michigan calls more than 500 households to conduct interviews. These interviews are conducted beginning either late in the previous month or early in the current month. Importantly, the final interviews are always completed before the end of the current month. So, in March for example, the first interview is conducted on February 26 and the last interview is completed by March 24. Importantly for our identification strategy, the CPI report for the reference month is always released the following month. For example, the March CPI report is released sometime in early April. Therefore, the assumption we make is that households' long-term inflation expectations can't respond within the month to the CPI release. After all, even a consumer that is eagerly awaiting the latest BLS report on price inflation in the current month won't be able to acquire this information until the following month. This argues for ordering CPI inflation after households' expectations for inflation over the next 5 to 10 years in a recursively identified VAR model.[19] Since we order energy and food inflation ahead of inflation expectations and CPI inflation, we interpret the orthogonalized innovation to CPI inflation as a core inflation shock.

4.3 The Forecast Consistent Prior

We specify our reduced form VAR as:

$$y_t = A_0 + A_1 y_{t-1} + A_2 y_{t-2} + A_3 y_{t-3} + u_t, u_t \sim \mathcal{N}(0, \Sigma),\ A = [A_0, A_1, A_2, A_3]' \tag{24}$$

We specify a non-informative prior, as described in Section 2.1, by setting V_0 to the zero matrix and $\nu_0 = 0$. Therefore, the posterior distributions for $\alpha = vec(A)$ and Σ, the VAR lag coefficients and the covariance matrix, are centered at the OLS estimates. Using our earlier parameterization of Σ from Section 2.1, $\Sigma = CQQ'C'$, where C is the lower-triangular Cholesky factor of Σ and Q is a square $m \times m$ orthogonal rotation matrix. To impose our recursive identification scheme motivated by the sequencing of surveys and data releases we set $Q = I$.

In addition to these dogmatic timing restrictions, we also employ our forecast consistent prior. In particular, we loosely impose some degree of consistency between the survey forecast and the VAR-implied forecast for the average CPI inflation rate over the 5 year period beginning 5 years from today. We specify our prior over the period h impulse response as follows:

$$\begin{aligned} g(A, C|h, \varepsilon_1^{\Delta p} = 1) = e'_{HH} IRF(A, C, h|\varepsilon_1^{\Delta p} = 1) \\ - \frac{1}{60} \sum_{j=1}^{60} \left[e'_{CPI} IRF(A, C, h + 60 + j|\varepsilon_1^{\Delta p} = 1) \right] \end{aligned} \tag{25}$$

where e_{HH} and e_{CPI} are selection vectors which respectively extract the responses of household inflation expectations and the CPI inflation rate. $IRF(A, C, h|\varepsilon_1^{\Delta p} = 1)$ denotes a vector with the impulse responses of all variables in period h to the structural innovation to $\varepsilon_1^{\Delta p}$ which denotes a price p (either energy, food, or core) inflation shock occurring in period 1.[20]

[19] Leduc, Sill and Stark (2007); Clark and Davig (2011); Leduc and Sill (2013) use a similar justification for their recursive identification schemes.

[20] We can stack the forecast consistency restriction in equation (25) for $h = 1, 2, \ldots, H$ to form the forecast consistent prior $g(A, C|1 : H, \varepsilon_1^{\Delta p} = 1) \sim \mathcal{N}(0_H, I_H \lambda^{-1})$ where λ tunes the precision over the forecast

Previous research has documented the potential for large discrepancies between household and VAR-based inflation forecasts. At a high level, the median of households' long-term inflation expectations are consistently higher than realized CPI inflation. Moreover, Coibion and Gorodnichenko (2012) and Coibion and Gorodnichenko (2015*a*) document that households' inflation forecast errors are predictable whereas, with the appropriate number of lags, VAR-based forecast errors are expected to be white noise. Unlike strict consistency, our forecast consistent prior can accommodate these discrepancies between household and VAR-based inflation forecasts, both on average and in response to specific structural shocks.

To calibrate λ, the degree of prior consistency imposed on the two forecasts, we use data from 1990-2003 as a training sample. While the impulse responses of inflation expectations to various inflationary shocks could change over different samples, the timing assumptions that drive our recursive identification strategy can be rationalized in this earlier sample. Therefore, we leverage these timing restrictions to identify energy, food, and core inflation shocks over the 1990-2003 sample. We then calibrate λ to match the observed discrepancy between the impulse responses of survey forecasts and VAR-implied forecasts over this training sample. Our calibration leads to relatively small values of λ for each of the three CPI inflation shocks. A low degree of forecast consistency is consistent with the large degree of information rigidities evident in households' forecasts of near-term inflation that Coibion and Gorodnichenko (2015*a*) document.

We simulate the posterior distribution of impulse responses by sequentially drawing from the posterior of $\Sigma = CC'$ and then $A|\Sigma$ M times. Each drawn pair of C and A implies an impulse response function for each of the three structural shocks we aim to identify. Therefore, for each draw of C and A and each structural shock $\varepsilon_1^{\Delta p} = 1$, we can calculate the posterior weight under the forecast consistent prior according to:

$$\begin{aligned}&w(A(m), C(m)|H, \varepsilon_1^{\Delta p} = 1) = \\ &\frac{exp(-0.5\lambda g(A(m), C(m)|H, \varepsilon_1^{\Delta p} = 1)' g(A(m), C(m)|H, \varepsilon_1^{\Delta p} = 1))}{\sum_{k=1}^{M} \exp(-0.5\lambda g(A(k), C(k)|H, \varepsilon_1^{\Delta p} = 1)' g(A(k), C(k)|H, \varepsilon_1^{\Delta p} = 1))}.\end{aligned}$$

These weights form the importance sampling weights we use to simulate the posterior set of impulse responses. We set $M = 5000$ and again scale the size of each draw to have the same initial effect on inflation before calculating the importance sampling weights.

consistent prior. In practice, we set H = 168 = 48+120 to form the prior over the 48 period impulse responses we plot.

4.4 Impulse Response Functions

We find evidence that households' longer-term inflation expectations became unresponsive to transitory inflation shocks after 2012. Figure 7 shows the responses of household's expected inflation over the next 5 to 10 years following energy inflation, food inflation, and core inflation shocks. The left column of Figure 7 shows the impulse responses over the 2004-2011 sample and the right column shows the impulse responses over the 2012-2019 sample.

The first row of Figure 7 shows that the response of household inflation expectations to energy inflation shocks has changed since 2012. Importantly, both samples include large booms and busts in oil prices — including 50 percent increases and declines over a 12-month window — which are thought to be a key factor influencing households' near-term inflation expectations (Coibion and Gorodnichenko, 2015*b*). The same also appears to be true for households' longer-term inflation expectations prior to 2012. However, after the 2012 adoption of a longer-run 2 percent inflation objective there is a marked change in the sensitivity of households' long-term inflation expectations to energy price shocks. The second row studies how the response of household inflation expectations to food inflation shocks has changed since 2012. While the spillovers to households' long-term inflation expectations were less pronounced prior to 2012, the complete lack of spillovers after 2012 is similar to what we observe for energy inflation shocks. Finally, the last row demonstrates that core inflation shocks also pass-through less to household's long-term inflation expectations since 2012. By the identifying assumptions, core inflation shocks have zero immediate pass-through to inflation expectations. But, prior to 2012, inflation expectations drifted higher with core inflation in subsequent periods. After 2012 there is no evidence of this delayed pass-through.

Our results in Figure 7 suggest that, prior to 2012, households' longer-run inflation expectations were influenced by changes in salient prices. However, we find evidence that households' longer-run inflation expectations no longer respond to such price changes after 2012. In contrast, the red-dashed lines in Figure 7 suggest that the VAR-implied forecasts for inflation over the next 5-10 years following each of these shocks is unresponsive in both samples. Consequently, restricting the VAR dynamics to impose a high-degree of forecast consistency would mask the recent changes in the dynamics of household inflation expectations. Therefore, this application highlights the flexibility of our flexible approach relative to strict consistency in a setting where the inflation expectations of the median University of Michigan survey respondent were historically governed by something other than our VAR model.

5 Inflation Tail Risks

According to former Federal Reserve Chairman, Alan Greenspan, "the conduct of monetary policy in the United States has come to involve, at its core, crucial elements of risk management." One element of this policy strategy involves managing risks around the Federal Reserve's price stability mandate. While during much of the 1980's and 1990's the FOMC was primarily concerned with defending its inflation mandate from above, in more recent decades it has confronted the risk of too low inflation. Concerns of deflation were especially elevated following the Global Financial Crisis and its aftermath, sparking broad interests in analyzing inflation tail risks over this period.

Several recent papers have measured inflation tail risks using financial market data and identify an elevated risk of deflation in the aftermath of the Global Financial Crisis. Fleckenstein, Longstaff and Lustig (2017) estimate a model of time-varying deflation risks identified from inflation swaps and options. Similarly, Anene and D'Amico (2017) and Hattori, Schrimpf and Sushko (2016) use inflation derivatives to study the impact of the FOMC's unconventional policy actions on stemming the risk of deflation. We contribute to the literature by studying inflation tail risks implied by a time-varying parameter (TVP-)VAR model of inflation that includes both near-term and longer-term survey forecasts of inflation. Relative to this previous literature, our approach has the appeal that, unlike financial market data, survey forecasts are not contaminated by time-varying risk premiums.[21]

This application illustrates several novel features of our forecast consistent prior that we have yet to demonstrate. First, in this application, the forecast consistent prior is applied in a setting with multiple survey measures at multiple horizons. Second, the setting for this application is a TVP-VAR model which highlights the computational ease of implementing the forecast consistent prior. Third, in this application, by simulating the predictive density of inflation from a VAR model with survey forecasts we demonstrate the appeal of augmenting VAR models with survey data to elicit moments of the forecast density that may not be explicitly provided by surveys. Finally, we apply the prior to unconditional forecasts from the VAR rather than impose restrictions on impulse response dynamics. Therefore, the tightness of the forecast consistent prior can be chosen to maximize the marginal likelihood.

[21] Our approach is comparable to Kozicki and Tinsley (2012) who estimate a time-varying parameter TVP-AR(p) model for inflation together with survey forecasts. However, they allow only constant terms ("inflation end-point") to drift over time while other coefficients in the AR(p) model remain time-invariant. Keeping AR coefficients time-invariant has some limitations in modeling changes in the relationship between inflation and inflation expectations For example, as our previous application showed, the adoption of explicit inflation target appear to have made inflation expectations less responsive to temporary shocks. For this reason, we consider a more general TVP-VAR model in which constant terms and lag coefficients drift over time.

Our results suggests that inflation expectations as measured from survey forecasts play an important role in limiting inflation tail risks. In particular, estimates from our forecast consistent TVP-VAR suggest that the risk of deflation during the Global Financial Crisis and its aftermath were generally lower than unconstrained VAR models and financial market estimates might suggest. The role of inflation expectations in influencing tail risks is especially vivid when survey forecasts for near-term inflation expectations hit all-time lows in 2009:Q1 which led deflation risks from the TVP-VAR model to sharply rise. However, as survey forecasts rebounded in subsequent quarters deflation probabilities quickly fell. In contrast, deflation risks from the TVP-VAR model without forecast consistency remain persistently elevated after 2009. The marginal likelihood criterion favors the TVP-VAR model with some degree of forecast consistency imposed, implying that the data prefers a model which tightly links deflation risks to variation in survey forecasts for inflation.

5.1 Data and TVP-VAR Model

The specification of our 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.[22] 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.

[22] The 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.

We model these five variables as a TVP-VAR(4) with stochastic volatility:

$$
\begin{aligned}
y_t &= A_{0,t} + \sum_{j=1}^{4} A_{j,t} y_{t-j} + u_t \,,\; u_t \sim \mathcal{N}(0, B^{-1}\Sigma_{u,t}B^{-1\prime}), \\
\Sigma_{u,t} &= \begin{pmatrix} \sigma_{1,t}^2 & 0 & \cdots & 0 \\ 0 & \sigma_{2,t}^2 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & \sigma_{5,t}^2 \end{pmatrix}, B = \begin{pmatrix} 1 & 0 & \cdots & 0 \\ B_{21} & 1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ B_{51} & \cdots & B_{54} & 1 \end{pmatrix}.
\end{aligned}
\tag{26}
$$

We detail the construction of our priors and the algorithm used to estimate and simulate the TVP-VAR model in the appendix.

5.2 The Forecast Consistency Prior

The TVP-VAR model contains both near-term and long-term survey forecasts of inflation as well as realized inflation. Therefore, we impose our forecast consistent prior over the survey and VAR-based forecasts for both forecast horizons. To ease the notation of writing out the forecast consistency restrictions, we express the TVP-VAR model in companion form:

$$
\tilde{y}_t = \tilde{A}_{0,t} + \tilde{A}_{1,t}\tilde{y}_{t-1} + \tilde{u}_t. \tag{27}
$$

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 the expectation operator E^e:

$$
\begin{aligned}
\pi_t^{e,L} &= \frac{\sum_{j=5}^{40} E_t^e(\pi_{t+j})}{36} \\
\pi_t^{e,S} &= \frac{\sum_{j=1}^{4} E_t^e(\pi_{t+j})}{4}
\end{aligned}
\tag{28}
$$

where E_t^e is the survey expectation when $e = S$ and the VAR-based expectation when $e = VAR$.[23]

[23] The end-point of the SPF 10-year inflation forecasts changes only in the first quarter of each year. Therefore, the number of quarterly inflation forecasts contained in the 10-year forecast varies depending on the quarter of year. We deal with this in the TVP-VAR by adjusting the lower limit of the sum and the denominator in π_t^L depending on the quarter of the year. Therefore, the notation below is illustrative and applies to the first quarter of the year.

Forecast consistency at both forecast horizons requires the following restrictions:

$$\begin{aligned}
\pi_t^{S,L} - \pi_t^{VAR,L} &= e'_{\pi,S,L}\tilde{y}_t - e'_\pi \frac{[\sum_{h=5}^{40}\sum_{j=0}^{h-1}\tilde{A}_{1,t}^j\tilde{A}_{0,t} + \sum_{h=5}^{40}\tilde{A}_{1,t}^h\tilde{y}_t]}{36}, \\
g(\tilde{A})_L &= [-e'_\pi \frac{[\sum_{h=5}^{40}\sum_{j=0}^{h-1}\tilde{A}_{1,t}^j\tilde{A}_{0,t}])}{36}, e'_{\pi,S,L} - e'_\pi \frac{[\sum_{h=5}^{40}\tilde{A}_{1,t}^h]}{36}]' \\
\pi_t^{S,S} - \pi_t^{VAR,S} &= e'_{\pi,S,S}\tilde{y}_t - e'_\pi \frac{[\sum_{h=1}^{4}\sum_{j=0}^{h-1}\tilde{A}_{1,t}^j\tilde{A}_{0,t} + \sum_{h=1}^{4}\tilde{A}_{1,t}^h\tilde{y}_t]}{4}, \\
g(\tilde{A})_S &= [-e'_\pi \frac{[\sum_{h=1}^{4}\sum_{j=0}^{h-1}\tilde{A}_{1,t}^j\tilde{A}_{0,t}]}{4}, e'_{\pi,S,S} - e'_\pi \frac{[\sum_{h=1}^{4}\tilde{A}_{1,t}^h]}{4}]', \\
g(\tilde{A}) &= [g(\tilde{A})_L, g(\tilde{A})_S]',
\end{aligned} \tag{29}$$

where e_i is a selection vector whose ith 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:

$$E_t^{\text{VAR}}(\prod_{k=1}^{h}\tilde{A}_{1,t+k}|\mathcal{F}_t) = \tilde{A}_{1,t}^h. \tag{30}$$

In this application we calibrate the hyperparamter λ that controls the tightness of forecast consistency prior restrictions by calculating the following marginal data density for different values of λ.[24]

$$p_(y^T|\lambda) = \int p(y^T|A^T, \Sigma_u^T, B)p(A^T, \Sigma_u^T, B)d(A^T, \Sigma_u^T, B). \tag{31}$$

The log marginal likelihood is maximized at $\lambda = 1.4$.[25] Therefore, imposing a modest degree of forecast consistency improves the time series fit of the TVP-VAR model.

5.3 Time-Varying Inflation Tail Risks

While survey forecasts are available for mean inflation outcomes, our interest in this application lies in assessing potential tail outcomes for inflation for which survey forecasts are not consistently available. Therefore, the forecast consistent prior is particularly useful as it tilts the mean of the predictive distribution of inflation in the TVP-VAR towards the available survey forecasts and then relies on the VAR model to generate other moments of the predictive distribution of inflation. We simulate the predictive distribution of inflation outcomes

[24] We calculate the inverse of the marginal likelihood using the harmonic mean of the likelihood implied by posterior draws. The calculations and details are explained in the appendix.

[25] 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.

using posterior draws of parameters and shocks up to time t from the TVP-VAR model. In practice, we achieve this by generating a full trajectory of inflation for the i-th posterior draw of parameters and shocks, denoted by $\pi_{t+h|t}(i)$, for each draw $i = 1, \cdots, M$ from the posterior. We then compute the probability that inflation will be less than 0 percent on average over the next H-periods according to:

$$P(\sum_{h=1}^{H} \pi_{t+h|t} < 0) = \frac{\sum_{i=1}^{M} \mathbb{I}(\sum_{h=1}^{H} \pi_{t+h|t}(i) < 0)}{M}. \tag{32}$$

We can perform a similar calculation to assess the likelihood of high inflation which, following Fleckenstein, Longstaff and Lustig (2017), we define as inflation above 4 percent.

Table 1 compares estimated deflation probabilities from the TVP-VAR with those from Fleckenstein, Longstaff and Lustig (2017) and the cross-sectional distribution of individual expectations from the University of Michigan consumer survey over the period of 2009:Q4-2015:Q4.[26] Although the underlying source data are completely different, the mean and median probabilities of deflation are quite comparable given the magnitude of the standard deviation of each measure. However, the TVP-VAR model with consistency restrictions implies uniformly lower deflation probabilities across all three quantile estimates (minimum, median, maximum) at both the 1- and 2-year horizons. This is also true for the estimated probability of high inflation (inflation greater than 4 percent) in Table 2.

The lower levels of inflation tail risks from our forecast consistent TVP-VAR model suggests that financial market measures overstate the tail risks to inflation. Moreover, simply including mean survey forecasts of inflation in the VAR fails to fully capture the potential role of survey expectations in driving inflation tail risks. The relative stability of professional forecasters' inflation expectations over the 2009-2015 sample reduces the overall threat of high inflation and deflation. More generally, the forecast consistency restrictions more tightly link inflation tail risks to survey forecasts of inflation. For example, in early 2009 as oil prices fell by more than $100 per barrel, survey forecasts for inflation over the next 1 year hit all time lows. The bottom row of Figure 8 shows that the risk of deflation sharply increased at this time. Then, as near-term inflation expectations rebounded, the risk of deflation subsided according to the forecast consistent TVP-VAR but remained elevated in the VAR model without the forecast consistency restrictions.

[26]For the University of Michigan survey, we calculate the percentage of respondents who anticipated prices would go down among all the respondents who provided answers on expected prices.

6 Conclusion

A growing literature has incorporated survey measures of expectations into VAR models to capture the importance of forward-looking behavior. In this paper, we have proposed the *forecast consistent prior* as a computationally efficient Bayesian framework for estimation and inference of VAR models with survey forecasts and realized data of a closely related variable. We highlight a range of possible applications of our framework in the context of both structural VAR shock identification as well as VAR forecasting. The applications shed light on the importance of synchronizing interest rate expectations to generate the desired effects from forward guidance, the degree to which households' longer-term inflation expectations became better anchored after the FOMC adopted an inflation target, and the role that professional forecasters' inflation expectations play in shaping inflation tail risks. However, many applications remain given the growing interest in identifying monetary and non-monetary news shocks as well as the increased interest in macroeconomics of understanding the formation and evolution of expectations. Therefore, as these literatures continue to advance, there will be a growing need for a framework such as ours that can efficiently and flexibly link VAR-based forecasts and survey forecasts without dogmatic restrictions.

References

Anene, Dominic, and Stefania D'Amico. 2017. "A Tale of Four Tails: Inflation, the Policy Rate, Longer-Term Rates, and Stock Prices." Federal Reserve Bank of Chicago Working Paper 2017-26.

Antolín-Díaz, Juan, and Juan F Rubio-Ramírez. 2018. "Narrative Sign Restrictions for SVARs." *American Economic Review*, 108(10): 2802–29.

Arias, Jonas E, Dario Caldara, and Juan F Rubio-Ramirez. 2018. "The systematic component of monetary policy in SVARs: an agnostic identification procedure." *Journal of Monetary Economics*. Forthcoming.

Barsky, Robert B., and Eric R. Sims. 2012. "Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence." *American Economic Review*, 102(4): 1343–77.

Barsky, Robert B, Susanto Basu, and Keyoung Lee. 2015. "Whither News Shocks?" *NBER Macroeconomics Annual*, 29(1): 225–264.

Binder, Carola. 2017. "Fed speak on main street: Central bank communication and household expectations." *Journal of Macroeconomics*, 52: 238–251.

Bundick, Brent, and A Lee Smith. 2018. "Does Communicating a Numerical Inflation Target Anchor Inflation Expectations? Evidence & Bond Market Implications." Federal Reserve Bank of Kansas City, Research Working Paper Working Paper 18-01.

Bundick, Brent, and A. Lee Smith. 2019. "The Dynamic Effects of Forward Guidance Shocks." *Review of Economics and Statistics*, Forthcoming(1).

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

Campbell, John Y, and Robert J Shiller. 1988. "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors." *The Review of Financial Studies*, 1(3): 195–228.

Cavallo, Alberto, Guillermo Cruces, and Ricardo Perez-Truglia. 2017. "Inflation expectations, learning, and supermarket prices: Evidence from survey experiments." *American Economic Journal: Macroeconomics*, 9(3): 1–35.

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.

Cogley, Timothy. 2005. "Changing Beliefs and the Term Structure of Interest Rates: Cross-equation Restrictions with Drifting Parameters." *Review of Economic Dynamics*, 8(2): 420–451.

Coibion, Olivier, and Yuriy Gorodnichenko. 2012. "What Can Survey Forecasts Tell Us about Information Rigidities?" *Journal of Political Economy*, 120(1): 116–159.

Coibion, Olivier, and Yuriy Gorodnichenko. 2015*a*. "Information rigidity and the expectations formation process: A simple framework and new facts." *American Economic Review*, 105(8): 2644–78.

Coibion, Olivier, and Yuriy Gorodnichenko. 2015*b*. "Is the Phillips curve alive and well after all? Inflation expectations and the missing disinflation." *American Economic Journal: Macroeconomics*, 7(1): 197–232.

Coibion, Olivier, Yuriy Gorodnichenko, and Saten Kumar. 2018. "How Do Firms Form Their Expectations? New Survey Evidence." *American Economic Review*, 108(9): 2671–2713.

Coibion, Olivier, Yuriy Gorodnichenko, Saten Kumar, and Mathieu Pedemonte. 2020. "Inflation expectations as a policy tool?" *Journal of International Economics*, 103297.

Cover, Thomas M, and Joy A Thomas. 2012. *Elements of information theory.* John Wiley & Sons.

Del Negro, Marco, and Frank Schorfheide. 2004. "Priors from General Equilibrium Models for VARs." *International Economic Review*, 45(2): 643–673.

Detmeister, Alan K, Daeus Jorento, Emily Massaro, and Ekaterina V Peneva. 2015. "Did the Fed's Announcement of an Inflation Objective Influence Expectations?" Board of Governors of the Federal Reserve System (US).

Doh, Taeyoung, and Amy Oksol. 2018. "Has the Anchoring of Inflation Expectations Changed in the United States during the Past Decade?" *Economic Review*, , (Q 1): 5–32.

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

Fleckenstein, Matthias, Francis A Longstaff, and Hanno Lustig. 2017. "Deflation risk." *The Review of Financial Studies*, 30(8): 2719–2760.

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

Gürkaynak, Refet S., Andrew Levin, and Eric Swanson. 2010. "Does Inflation Targeting Anchor Long-Run Inflation Expectations? Evidence from the U.S., U.K, and Sweden." *Journal of the European Economic Association*, 8(6): 1208–1242.

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*.

Hattori, Masazumi, Andreas Schrimpf, and Vladyslav Sushko. 2016. "The Response of Tail Risk Perceptions to Unconventional Monetary Policy." *American Economic Journal: Macroeconomics*, 8(2): 111–36.

Keating, John W. 1990. "Identifying VAR Models Under Rational Expectations." *Journal of Monetary Economics*, 25(3): 453–476.

Kilian, Lutz, and Helmut Lütkepohl. 2017. *Structural vector autoregressive analysis.* Cambridge University Press.

Kilian, Lutz, and Logan T Lewis. 2011. "Does the Fed Respond to Oil Price Shocks?" *The Economic Journal*, 121(555): 1047–1072.

Kim, Jae-Young. 2002. "Limited Information Likelihood and Bayesian Analysis." *Journal of Econometrics*, 107(1-2): 175–193.

Kozicki, Sharon, and Peter A Tinsley. 2012. "Effective Use of Survey Information in Estimating the Evolution of Expected Inflation." *Journal of Money, Credit and Banking*, 44(1): 145–169.

Leduc, Sylvain, and Keith Sill. 2013. "Expectations and Economic Fluctuations: An Analysis Using Survey Data." *Review of Economics and Statistics*, 95(4): 1352–1367.

Leduc, Sylvain, Keith Sill, and Tom Stark. 2007. "Self-Fulfilling Expectations and the Inflation of the 1970s: Evidence from the Livingston Survey." *Journal of Monetary economics*, 54(2): 433–459.

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

Melosi, Leonardo. 2017. "Signalling effects of monetary policy." *The Review of Economic Studies*, 84(2): 853–884.

Moon, Hyungsik Roger, and Frank Schorfheide. 2012. "Bayesian and frequentist inference in partially identified models." *Econometrica*, 80(2): 755–782.

Nakamura, Emi, and Jón Steinsson. 2018. "High-frequency identification of monetary non-neutrality: the information effect." *The Quarterly Journal of Economics*, 133(3): 1283–1330.

Ramey, Valerie A. 2016. "Macroeconomic shocks and their propagation." In *Handbook of Macroeconomics.* Vol. 2, 71–162. Elsevier.

Robert, Christian. 2007. *The Bayesian Choice.* Springer Science & Business Media.

Sims, Eric R. 2012. "News, non-invertibility, and structural VARs." *Advances in Econometrics*, 28: 81.

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.

Uhlig, Harald. 2005. "What are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure." *Journal of Monetary Economics*, 52(2): 381–419.

Uhlig, Harald. 2017. "Shocks, Sign Restrictions, and Identification." Vol. 59, 95, Cambridge University Press.

Wolf, Christian K. 2018. "SVAR (mis-)identification and the real effects of monetary policy."

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

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.

Yellen, Janet L. 2016. "The Federal Reserve's Monetary Policy Toolkit: Past, Present and Future. Speech given at the Federal Reserve Bank of Kansas City Economic Symposium, Jackson Hole, Wyoming, August 26, 2016." Board of Governors of the Federal Reserve System (U.S.) Speech.

Table 1: Summary Statistics for Deflation Probabilities: 2009:Q4 - 2015:Q4

Horizon	Source	Mean	Standard Deviation	Minimum	Median	Maximum
1 year	Fleckenstein et al. (2017)	18.759	10.155	2.003	15.739	47.757
1 year	University of Michigan Survey	15.906	6.1249	6.06	14.1414	32.6531
1 year	TVP-VAR($\lambda = 0$)	13.74	9.62	2.72	10.54	45.98
1 year	TVP-VAR($\lambda = 1.4$)	9.03	8.22	1.66	7.32	42.78
2 years	Fleckenstein et al. (2017)	13.799	7.261	1.801	11.621	35.421
2 years	TVP-VAR ($\lambda = 0$)	16.34	6.27	8.38	15.14	32.06
2 years	TVP-VAR ($\lambda = 1.4$)	9.54	4.59	4.1	7.98	23.96

Notes: Deflation probabilities from Fleckenstein et al. (2017) are based on daily observations on inflation swaps and options from October 5, 2009 to October 28, 2015 while those from the University of Michigan survey are quarterly average values of monthly observations from October, 2009 to October, 2015.

Table 2: Summary Stastistics for High Inflation (>4 percent) Probabilities: 2009:Q4 - 2015:Q4

Horizon	Source	Mean	Standard Deviation	Minimum	Median	Maximum
1 year	Fleckenstein et al. (2017)	2.615	2.46	0.139	2.086	16.268
1 year	TVP-VAR ($\lambda = 0$)	1.34	2.3	0.04	0.4	10.92
1 year	TVP-VAR ($\lambda = 1.4$)	0.64	1.29	0	0.16	6.06
2 years	Fleckenstein et al. (2017)	3.089	2.394	0.295	2.702	15.315
2 years	TVP-VAR ($\lambda = 0$)	4.10	2.86	0.88	3.14	14.02
2 years	TVP-VAR ($\lambda = 1.4$)	2.3	2.17	0.32	1.32	9.64

Notes: Deflation probabilities from Fleckenstein et al. (2017) are based on daily observations on inflation swaps and options from October 5, 2009 to October 28, 2015.

Figure 1: Refining Sign Restrictions with the Forecast Consistent Prior

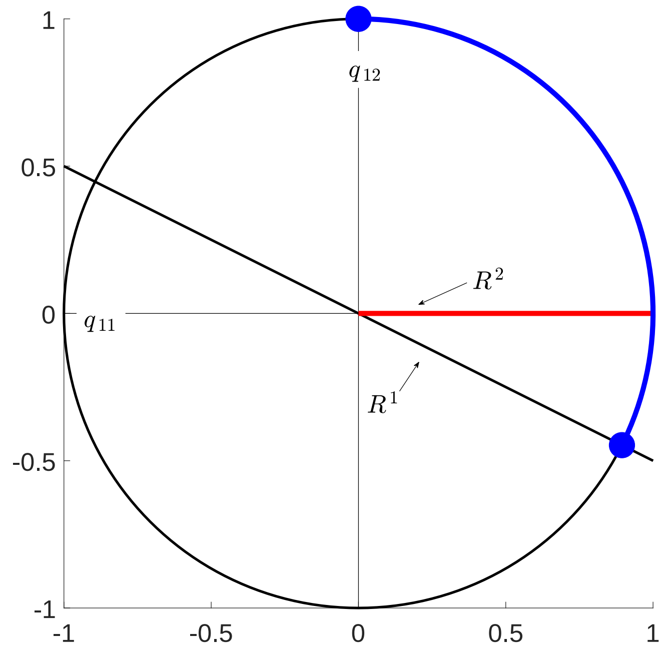

(a) Identified Set: Normalization and Sign Restriction

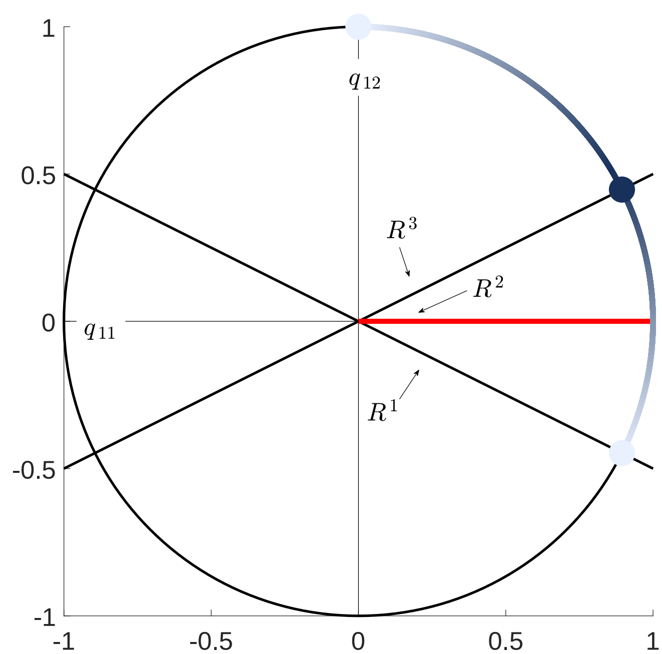

(b) Identified Set: Normalization, Sign Restriction, and Forecast Consistency

Figure 2: Forward Guidance Shock: Sign Restrictions Only

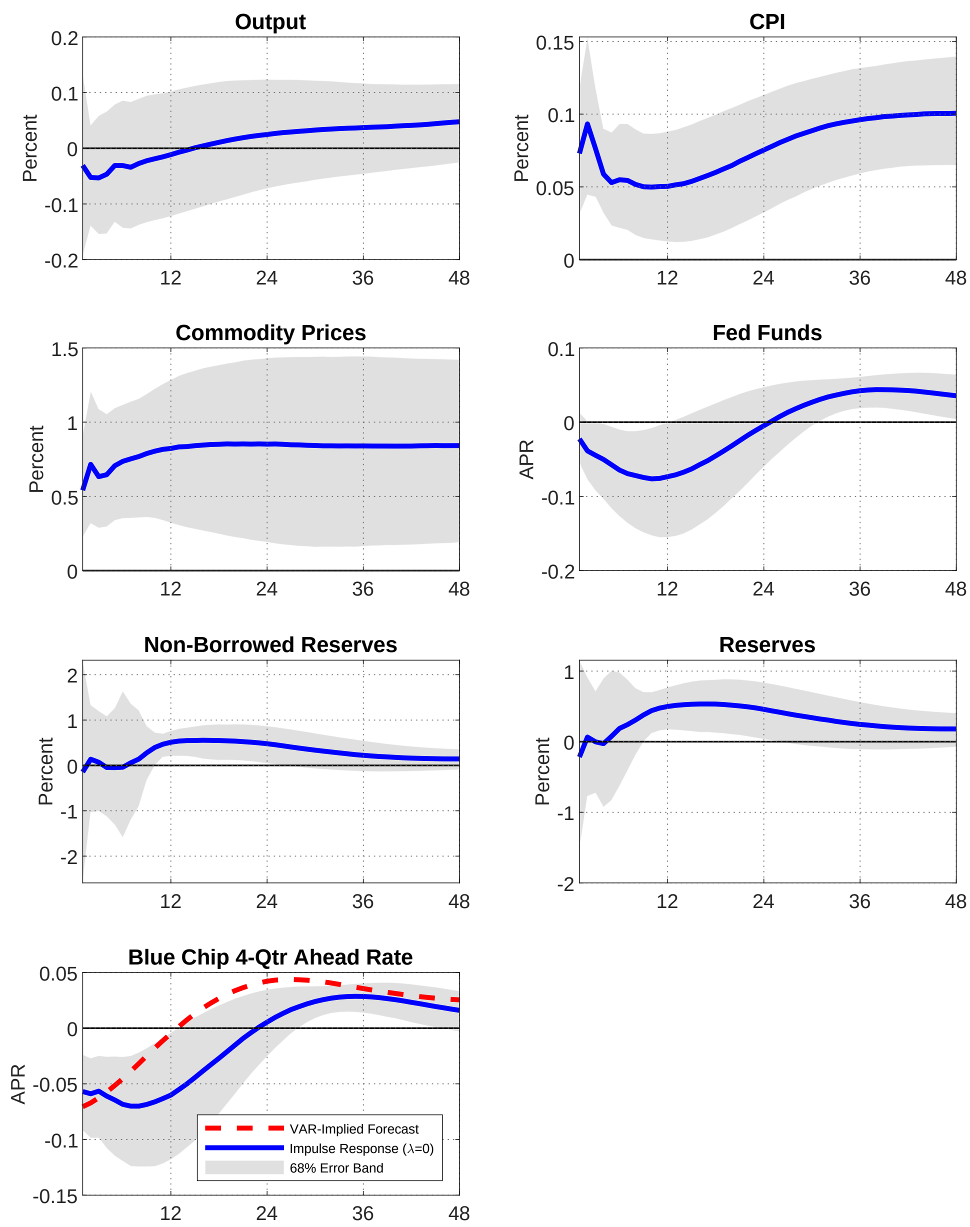

Notes: This figure shows the impulse responses to an identified forward guidance shock using only sign restrictions. The solid blue line is the median response and the shaded region is the 68% interval among structural VAR models. The red-dashed line shows the VAR-implied response of future short-term interest rates. The estimation sample period is 1994-2007.

Figure 3: Forward Guidance Shocks: "Best" and "Worst" Fitting Draws

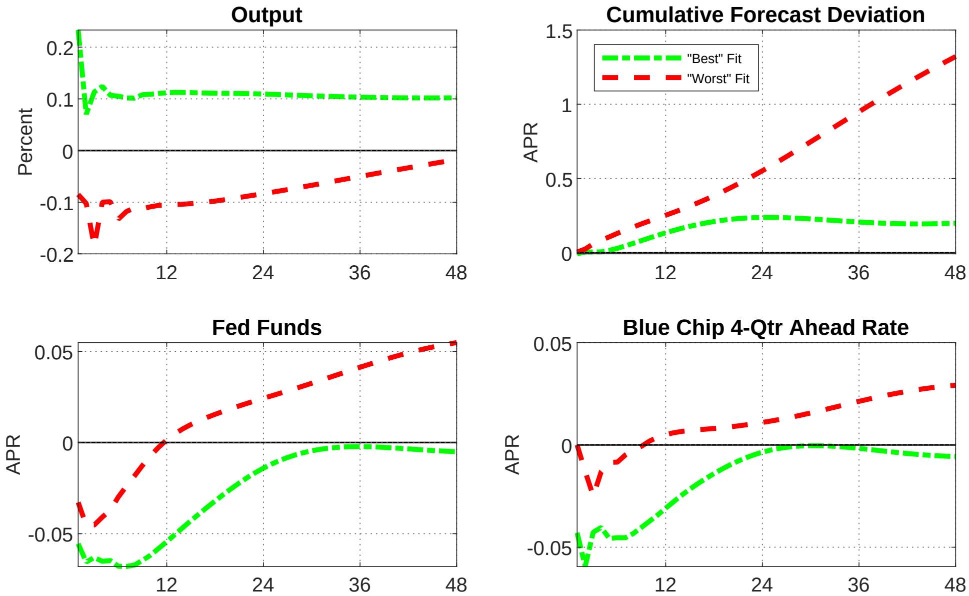

Notes: This figure shows the impulse responses to an identified forward guidance shock for the "best" and "worst" fitting models. The "best" fitting model is the draw that comes the closest to satisfying the forecast consistency restrictions and the "worst" fitting model is the draw that is the furthest from satisfying the forecast consistency restrictions. The estimation sample period is 1994-2007.

Figure 4: Output Effects of Forward Guidance: The Role of the Forecast Consistent Prior

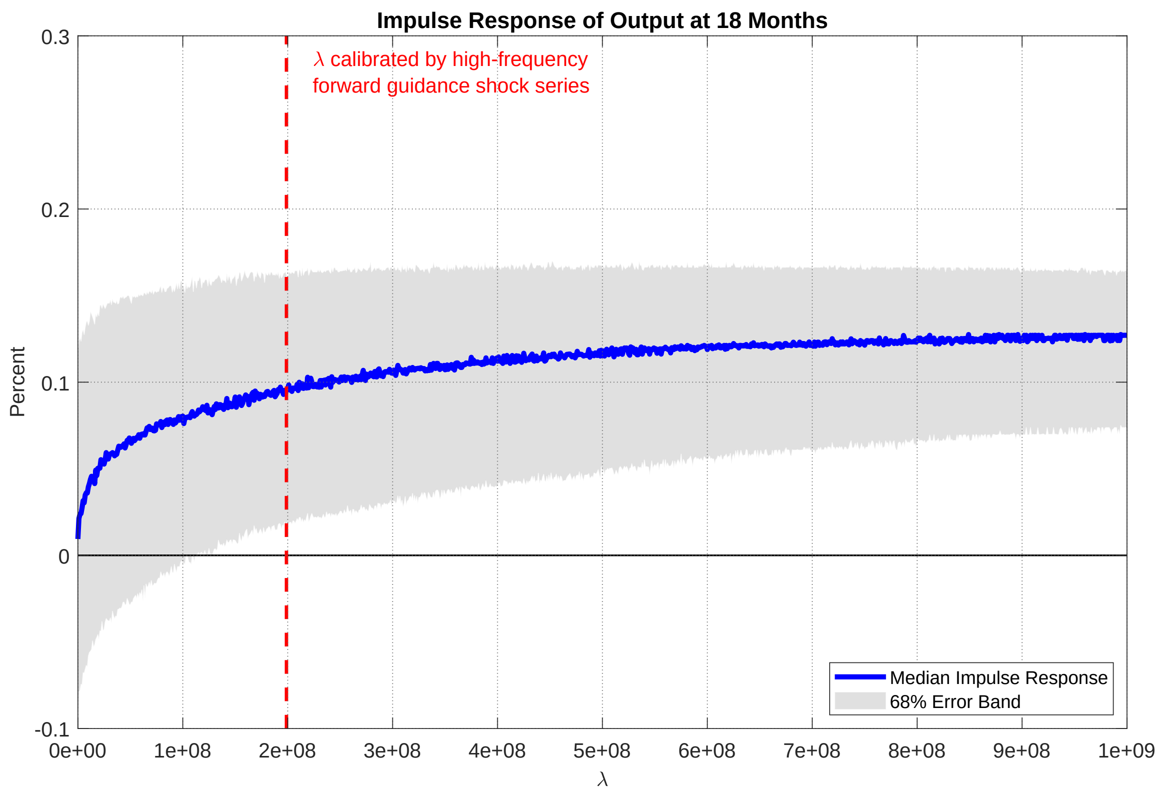

Notes: This figure shows the median output response and the corresponding 68% intervals after 18 months for alternative values of λ, which governs the tightness of the forecast consistent prior. The vertical red line denotes the value of λ selected to maximize the correlation between the structural VAR series of forward guidance innovations and high-frequency financial market measures of forward guidance shocks. The estimation sample period is 1994-2007.

Figure 5: Forward Guidance Shock: Sign Restrictions and the Forecast Consistent Prior

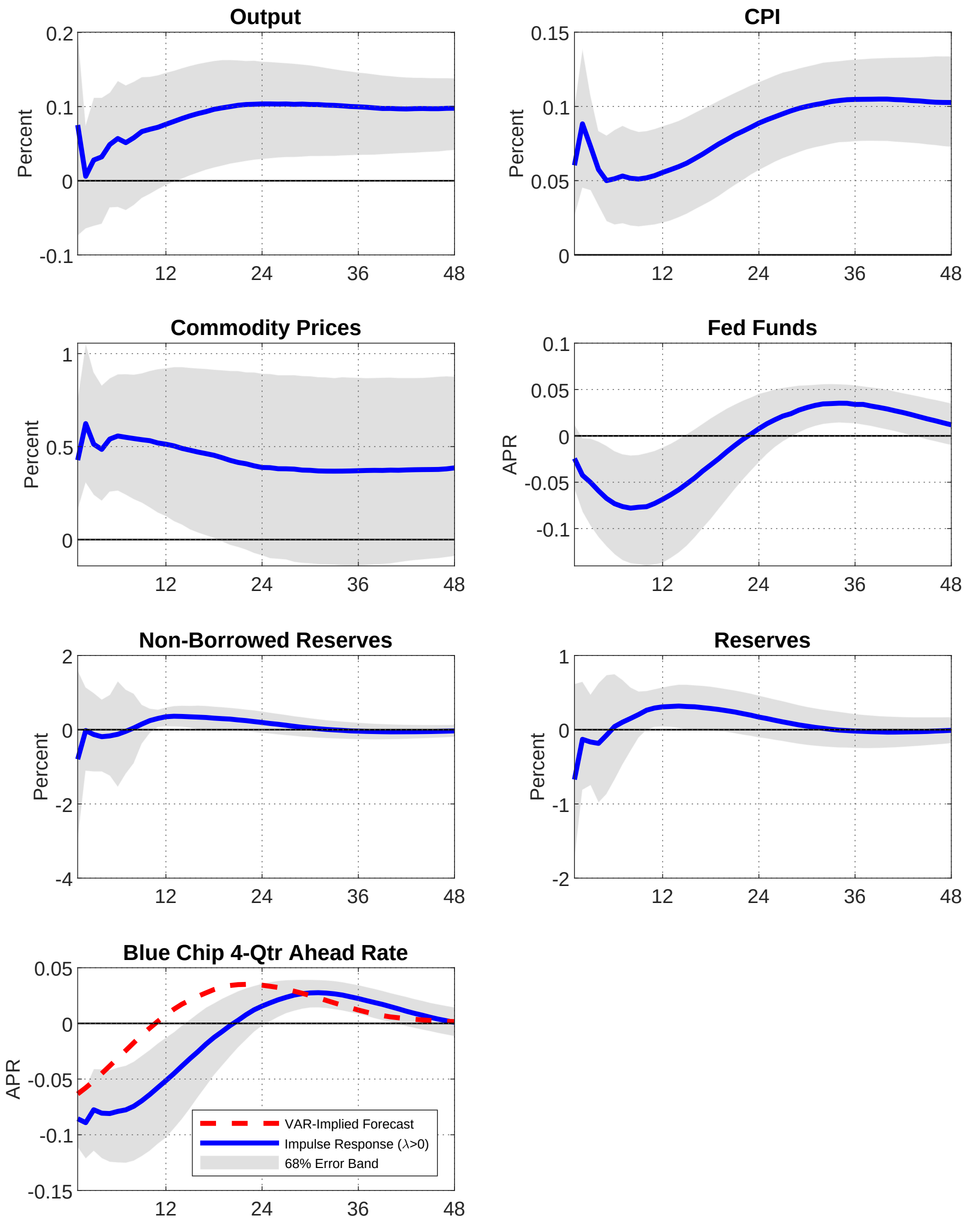

Notes: This figure shows the impulse responses to an identified forward guidance shock using sign restrictions as well as our forecast consistent prior. The solid blue line is the median response and the shaded region is the 68% error band. The red line shows the VAR-implied response of future short-term interest rates. The estimation sample period is 1994-2007.

Figure 6: Forward Guidance Shock: Alternative Specifications

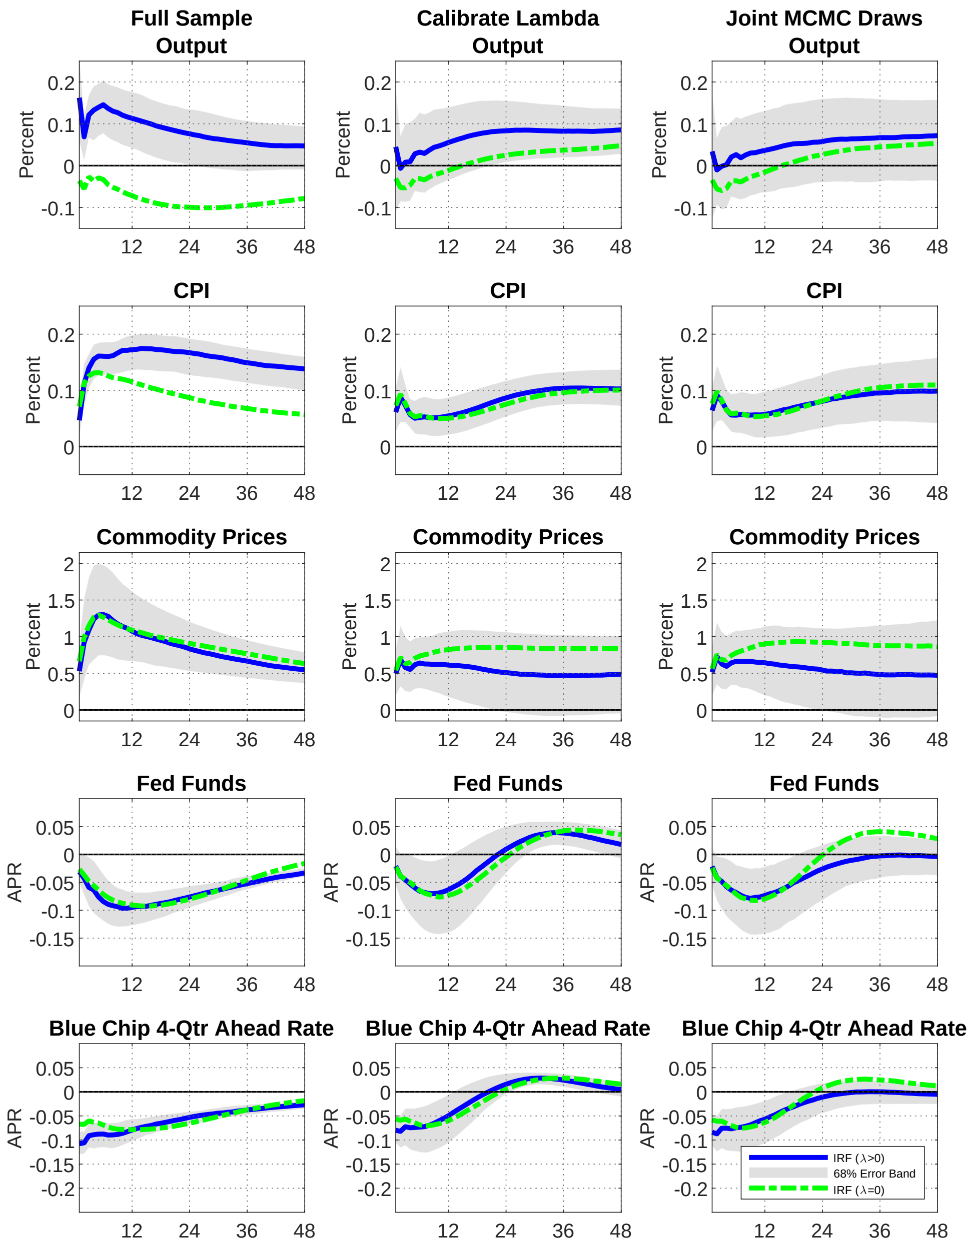

Notes: This figure shows the impulse responses to an identified forward guidance shock using sign restrictions as well as our forecast consistent prior. The solid blue line is the median response and the shaded region is the 68% error band. The green-dashed line shows the median impulse response to an identified forward guidance shock using only sign restrictions. Each column shows impulse responses from an alternative VAR model. For details, see Section 3.5.

Figure 7: Anchoring of Consumer Inflation Expectations

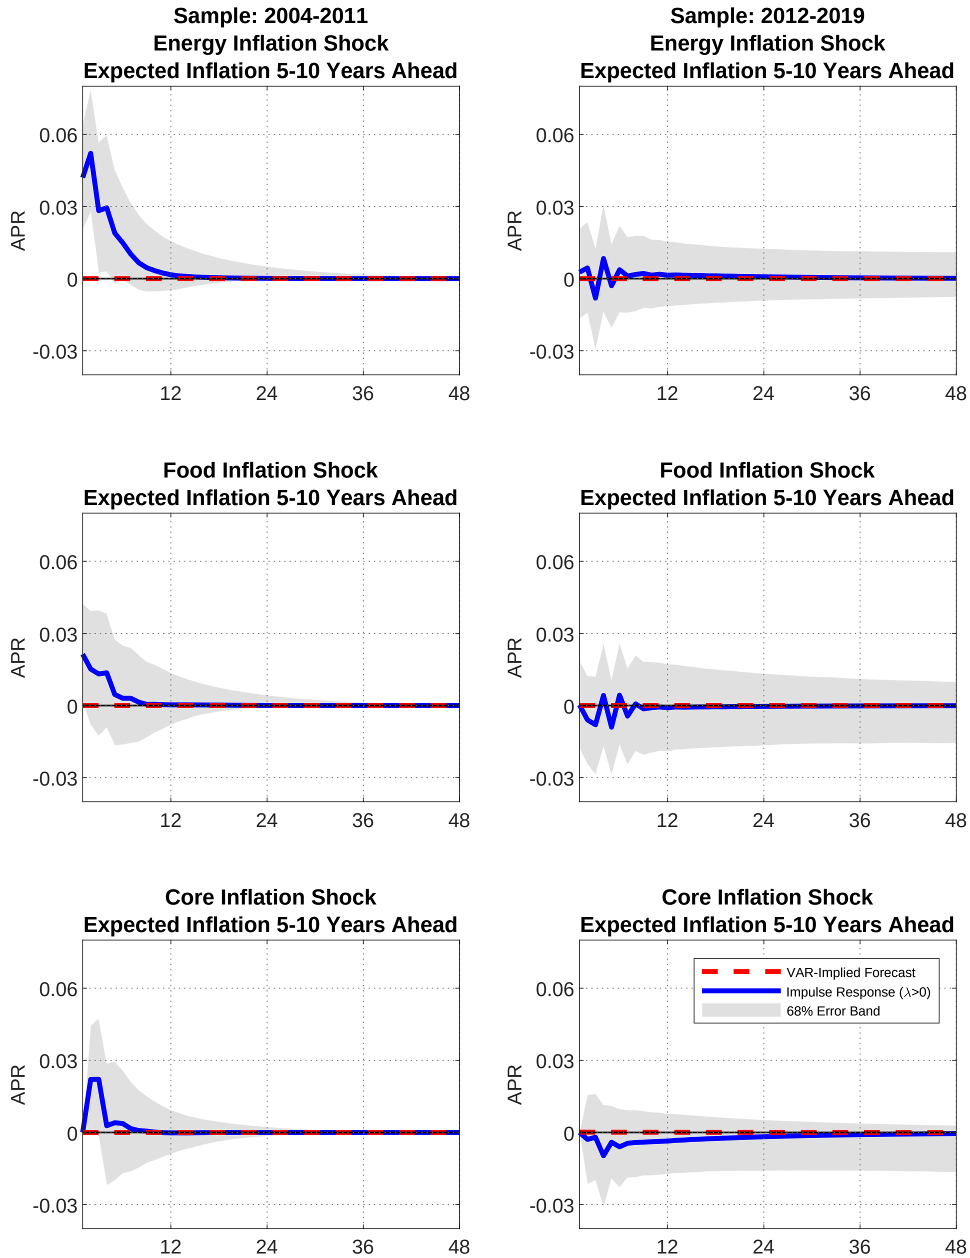

Notes: This figure shows the impulse responses of households' expectations for inflation over the next 5 to 10 years following an identified shock to energy inflation, food inflation, and core inflation. The solid blue line is the median response and the shaded region is the 68% error band for the University of Michigan's Survey of Consumers. The red-dashed line shows the median impulse response of the VAR-implied forecast of inflation over the next 5 to 10 years.

Figure 8: Inflation Tail Risks

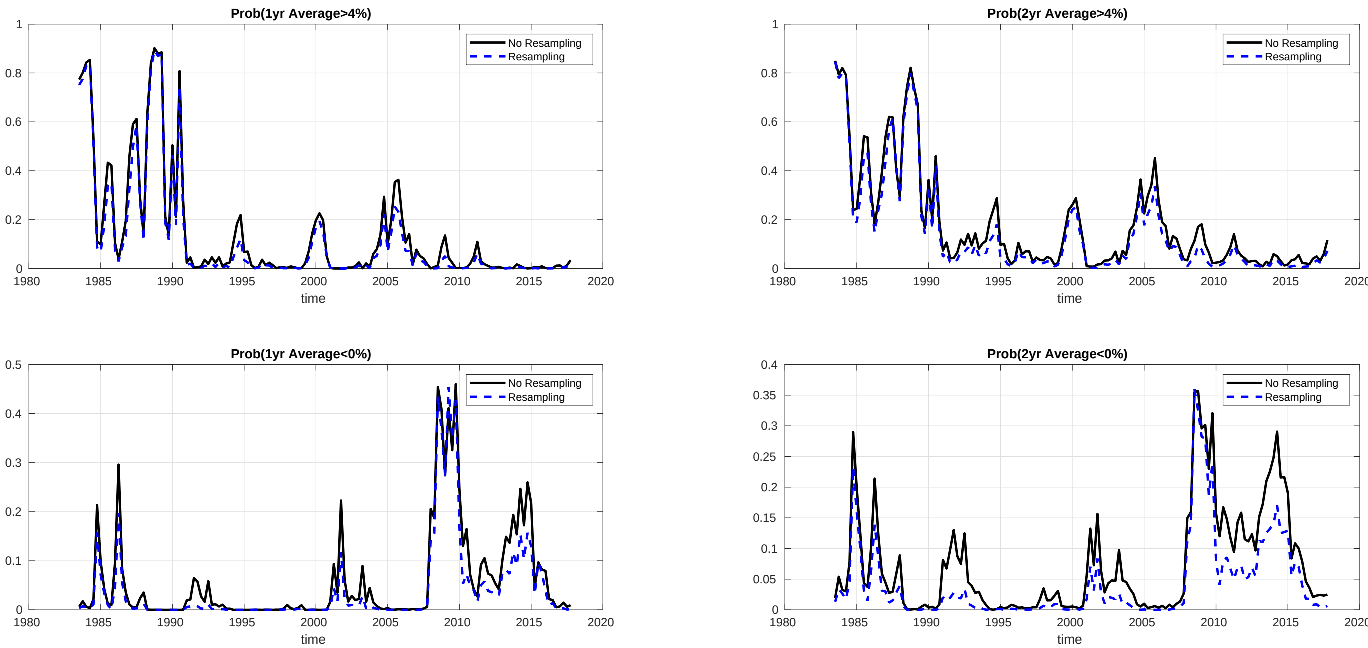

Notes: This figure shows the time-varying probabilities of high inflation (inflation >4 percent) in the top row and the time-varying probability of deflation (inflation <0 percent) from our TVP-VAR model. The left column shows these probabilities over the 1 year horizon and the right column shows these probabilities over the 2 year horizon. The solid black lines show these probabilities from the TVP-VAR without the forecast consistent prior, labeled "No Resampling", which corresponds to $\lambda = 0$ and the dotted blue lines show these probabilities from the TVP-VAR with the forecast consistent prior, labeled "Resampling", which corresponds to $\lambda > 0$.