Structural Scenario Analysis with SVARs[1]

Juan Antolín-Díaz[a], Ivan Petrella[b], Juan F. Rubio-Ramírez[c]

[a] *London Business School*
[b] *University of Warwick*
[c] *Emory University, Federal Reserve Bank of Atlanta*

Abstract

Macroeconomists seeking to construct conditional forecasts often face a choice between taking a stand on the details of a fully-specified structural model or relying on empirical correlations from vector autoregressions and remain silent about the underlying causal mechanisms. This paper develops tools for constructing "structural scenarios" that can be given an economic interpretation using identified structural VARs. We provide a unified and transparent treatment of conditional forecasting and structural scenario analysis and relate our approach to entropic forecast tilting. We advocate for a careful treatment of uncertainty, making the methods suitable for density forecasting and risk assessment. We also propose a metric to assess and compare the plausibility of alternative scenarios. We illustrate our methods with two applications: assessing the power of forward guidance about future interest rate policies and stress testing the reaction of bank profitability to an economic recession.

Keywords: Conditional forecasts, SVARs, Bayesian methods, forward guidance, stress testing.

1. Introduction

Macroeconomists and policy analysts are routinely required to answer questions of the form: "If, for the next few quarters, variable x were to follow alternative paths, how would

[1]We are grateful to Nabeel Abdoula, Piotr Chmielowski, Gavyn Davies, Domenico Giannone, Jonas Hallgren, Marek Jarocinski, Giulio Nicoletti, Giuseppe Ragusa, Paolo Surico, Fabrizio Venditti, and seminar participants at Warwick Business School, European Central Bank, Bank of England, London Business School, Bank of Spain, Banca d'Italia, and Central Bank of Ireland for helpful comments and suggestions. Yad Selvakumar provided excellent research assistance. The views expressed here are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Atlanta or the Federal Reserve System.

the forecasts of other variables change?" Uses of these conditional forecasts include assessing the impact of alternative paths for policy instruments, incorporating external information into model forecasts, and "stress testing" asset prices or bank profits against the onset of an economic downturn. In practice, researchers often face a choice between taking a stand on the details of a fully-fledged DSGE model (see Del Negro and Schorfheide, 2013) or relying on reduced-form conditional forecasts from vector autoregression (VAR) models, as proposed by Doan et al. (1986) and Waggoner and Zha (1999). A limitation of the latter approach is that it relies on the empirical correlations among the variables in the system, and is therefore silent about the underlying causal mechanisms behind the results. This can often severely confound the interpretation of conditional forecasts.

This paper develops a unified and transparent treatment of both conditional forecasting and structural scenario analysis with structural VAR (SVAR) models. We define a structural scenario as a combination of a path for one or more variables in the system, and a restriction that only a chosen subset of structural shocks can deviate from their unconditional distribution. This exercise requires identifying economically meaningful shocks. We provide closed-form solutions for the distribution of the variables and the underlying shocks under the restrictions. In addition, we demonstrate the equivalence of our formulae to the entropic tilting method popularized by Robertson et al. (2005). Another contribution, which follows naturally from the connection to entropic tilting, is to propose a metric based on the Kullback-Leibler divergence to assess the plausibility of a given scenario. Generalizing the "modest policy interventions" of Leeper and Zha (2003), it quantifies the degree to which a scenario involves a moderate deviation from the unconditional distribution of the shocks, in which case the results of the analysis are likely to be reliable. In this way, we follow the tradition of Sims (1986) in wanting to provide practical tools for macroeconomic analysis whenever one does not want to commit to a particular DSGE model, while retaining a healthy dose of skepticism about the limits of structural scenarios using linear models, as stressed by Lucas (1976).

Illustrative example. To highlight the importance of specifying the structural shocks driving a particular scenario, consider an SVAR including output growth, inflation, and the fed funds rate, which will be analyzed in detail in our empirical application. Figure 1 compares the

different answers that arise from the classic conditional forecasting question ("What is the likely path of output and inflation, given that the fed funds rate is reduced to 1% and kept there for two years?") and an alternative one, which we call a structural scenario: "What is the likely path of output and inflation, if monetary policy shocks lower the fed funds rate to 1% and keep it there for two years?"[2] In the latter case, only the policy shocks are allowed to deviate from their unconditional distribution to match the desired interest rate path. As we can see, for the exact same path of the interest rate, inflation and output are lower than in the unconditional forecast in the first exercise, but higher in the second one.

The differences are not surprising once we recognize that the bulk of the movements in the fed funds rate represents the systematic reaction of the Fed to output and inflation developments (see Leeper et al., 1996). The classic conditional forecasting question is really asking: "What is the most likely set of circumstances under which the Fed might keep the fed funds rate at zero for two years?" The answer is "subdued output growth and inflation." If communicated to the public as forward guidance, this is the kind of exercise that Campbell et al. (2012) call "Delphic": it provides a forecast of macroeconomic outcomes and likely policy actions based on the policymaker's potentially superior information about future economic fundamentals. The structural scenario exercise is instead "Odyssean" in the sense that it does not reveal any information about future fundamentals, but rather can be interpreted as an intention to keep the policy rate constant whatever the fundamentals happen to be. Clearly, in order to produce the latter we will need to identify, at least, the monetary policy shock, and the results will depend on the identification scheme used. On the contrary, identification was not necessary for the Delphic exercise.

Relation to the literature. Classic conditional forecasts and counterfactual analyses with VARs have a long tradition in macroeconomics, going back at least to Doan et al. (1986). Working in the reduced-form setting, Clarida and Coyle (1984) propose an implementation based on the Kalman filter, whereas Waggoner and Zha (1999) provide well-known methods for computing conditional density forecasts with VAR models. Their results have been extended to large

[2]Since 2009 the FOMC communicates interest rate targets as a range, e.g., "1-1.25%". For brevity, we always refer to the lower end in the text, but we implement all our empirical scenarios using the mid-point.

Figure 1: Answers to Two Alternative Questions

(a) Conditional Forecast ('Most likely path for output and inflation given path for FF rate')

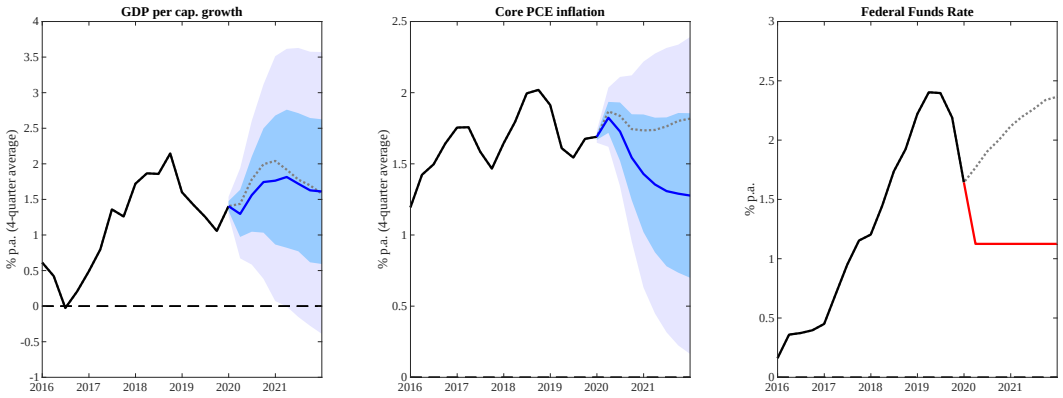

(b) Structural Scenario ('Most likely path for output and inflation if monetary policy shocks drive path for FF rate')

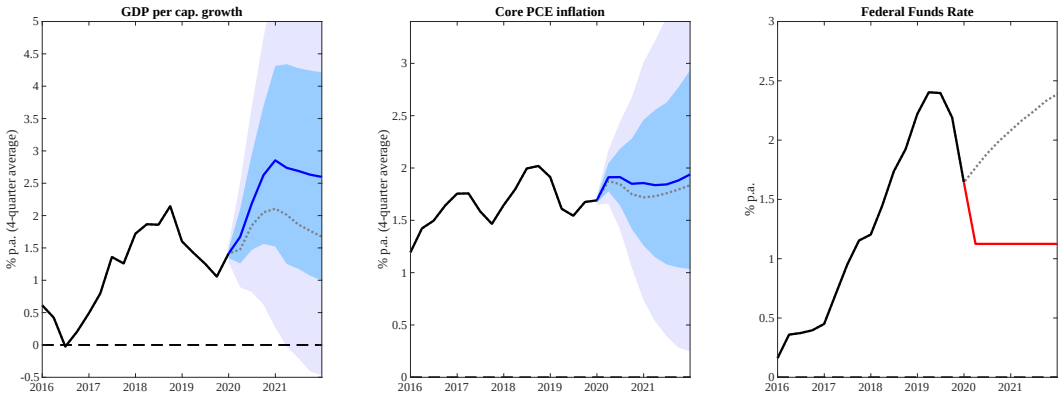

Note: For each panel, the solid black lines represent actual data, the solid red line is the conditioning assumption on the fed funds rate, the solid blue line is the median forecast for the remaining variables and periods, and the blue shaded areas denote the 40 and 68% pointwise credible sets around the forecasts. The dotted gray lines represent the median unconditional forecast.

systems (Banbura et al., 2015) and uncertainty about the paths of the conditioning variables (Andersson et al., 2010). The use of conditional forecasts has increased dramatically after the 2008-09 financial crisis, with a variety of applications including assessing structural change (Aastveit et al., 2017; Giannone et al., 2019), evaluating of non-standard monetary policies (Giannone et al., 2012; Altavilla et al., 2016), nowcasting and forecasting (Giannone et al., 2014; Tallman and Zaman, 2018), and evaluating of macro-prudential policies (Conti et al., 2018). We stress that all of this recent wave of research is reduced-form in nature, conditioning

exclusively on the path of observable variables, even when a structural interpretation is at times intended. Our suggestion is that scenarios should be instead constructed on the basis of economically meaningful shocks.

There is also a strand of the literature that works with identified SVARs to analyze monetary policy alternatives (Leeper and Zha, 2003) or the impact of oil supply shocks (Baumeister and Kilian, 2014). In a similar spirit, counterfactual analysis exercises were carried out by Bernanke et al. (1997), Hamilton and Herrera (2004) and Kilian and Lewis (2011), among others. Relative to existing approaches, our proposal carries several innovations. First, we provide closed-form solutions that are valid for an arbitrary number of conditions. Second, our definition of scenario avoids the need to elicit a priori conditions on the values of (unobserved) structural shocks. Instead, we specify conditions on observable variables, and select which shock(s) are driving the forecasts. We believe our approach is much more intuitive and will greatly facilitate practical application. Finally, we provide for a more complete treatment of uncertainty. Specifically, our methods allow us to recognize the uncertainty about the values of the "non-driving" shocks as well as the conditioning paths for the variables themselves, which the existing approaches have tended to shut down.[3] As we will see in our empirical applications, ignoring these sources of uncertainty may lead to excessively narrow confidence bands around forecasts. The increasing usage of density forecasts and fan charts for policy communication and macreoconomic analysis, as well as the recent emphasis on macroeconomic risk management (see Kilian and Manganelli, 2008), makes the treatment of uncertainty a first-order concern.

Organization of the paper. The rest of this paper is organized as follows. Section 2 presents the general econometric framework and formalizes the concept of structural scenarios. Section 3 introduces a measure of the plausibility of the different scenarios. Section 4 describes

[3]Additionally, by extending the existing Bayesian methods to set and partially identified SVARs, we also consider uncertainty originating both from the estimation of the reduced-form parameters and from the identification restrictions. The existing procedures, such as Baumeister and Kilian (2014) and Clark and McCracken (2014), usually ignored parameter uncertainty. Waggoner and Zha (1999) stress that ignoring parameter uncertainty can potentially result in misleading conditional forecasts. Of course, our methods can be applied as well to exactly and fully identified SVARs, where the latter source of uncertainty is not a concern.

algorithms to implement our techniques. Sections 5 and 6 illustrate our techniques using two examples. First, we further develop the monetary example above and explore the effects of forward guidance and average inflation targeting. Second, we consider a larger SVAR with macro and financial variables, and carry out a "stress testing" exercise to assess the response of asset prices and bank profitability to an economic recession.

2. Econometric framework

We work with the SVAR, written compactly

$$\mathbf{y}_t'\mathbf{A}_0 = \mathbf{x}_t'\mathbf{A}_+ + \boldsymbol{\varepsilon}_t' \text{ for } 1 \leq t \leq T, \tag{1}$$

with $\mathbf{A}_+' = \begin{bmatrix}\mathbf{A}_1' & \cdots & \mathbf{A}_p' & \mathbf{d}'\end{bmatrix}$, $\mathbf{x}_t' = \begin{bmatrix}\mathbf{y}_{t-1}', \ldots, \mathbf{y}_{t-p}', 1\end{bmatrix}$, where $\mathbf{y}_t$ is an $n \times 1$ vector of observables, $\boldsymbol{\varepsilon}_t$ is an $n \times 1$ vector of structural shocks, $\mathbf{A}_\ell$ is an $n \times n$ matrix of parameters for $0 \leq \ell \leq p$ with $\mathbf{A}_0$ invertible, $\mathbf{d}$ is an $1 \times n$ vector of parameters, p is the lag length, and T is the sample size. The vector of shocks $\boldsymbol{\varepsilon}_t$, conditional on past information and the initial conditions $\mathbf{y}_0, ..., \mathbf{y}_{1-p}$, is distributed $\mathcal{N}(\mathbf{0}_{n\times 1}, \mathbf{I}_n)$, where $\mathbf{0}_{n\times 1}$ is an $n \times 1$ matrix of zeros and $\mathbf{I}_n$ is an $n \times n$ identity matrix. The reduced-form representation implied by Equation (1) is

$$\mathbf{y}_t' = \mathbf{x}_t'\mathbf{B} + \mathbf{u}_t' \text{ for } 1 \leq t \leq T, \tag{2}$$

where $\mathbf{B} = \mathbf{A}_+\mathbf{A}_0^{-1}$, $\mathbf{u}_t' = \boldsymbol{\varepsilon}_t'\mathbf{A}_0^{-1}$, and $\mathbb{E}[\mathbf{u}_t\mathbf{u}_t'] = \boldsymbol{\Sigma} = (\mathbf{A}_0\mathbf{A}_0')^{-1}$. The matrices $\mathbf{B}$ and $\boldsymbol{\Sigma}$ are the reduced-form parameters, while $\mathbf{A}_0$ and $\mathbf{A}_+$ are the structural parameters. Similarly, $\mathbf{u}_t'$ are the reduced-form innovations. While the shocks are orthogonal and have an economic interpretation, the innovations may be correlated and do not have an interpretation.

As is well known, the model defined in Equation (1) has an identification problem. As described in Rubio-Ramirez et al. (2010), one can reparameterize this model in terms of $\mathbf{B}$ and $\boldsymbol{\Sigma}$ together with an $n \times n$ orthogonal rotation matrix $\mathbf{Q}$, such that for a given $\mathbf{B}$ and $\boldsymbol{\Sigma}$, a choice of $\mathbf{Q}$ implies a particular, observationally equivalent choice of structural parameters. To solve the identification problem, one often imposes restrictions on either the structural parameters or some function–such as the impulse response functions (IRFs)–thereof, pinning

down a particular $\mathbf{Q}$ (point identification) or narrowing down a set of $\mathbf{Q}$'s (set identification).

2.1. Unconditional forecasting

Assume that we want to forecast the observables for h periods ahead using the VAR in Equations (1)-(2). Given the history of observables $\mathbf{y}^T = \left(\mathbf{y}'_{1-p} \dots \mathbf{y}'_T\right)'$, the unconditional forecast of the observables, i.e., $\mathbf{y}'_{T+1,T+h} = \left(\mathbf{y}'_{T+1} \dots \mathbf{y}'_{T+h}\right)$ can be written

$$\mathbf{y}_{T+1,T+h} = \mathbf{b}_{T+1,T+h} + \mathbf{M}'\boldsymbol{\varepsilon}_{T+1,T+h}, \tag{3}$$

where $\boldsymbol{\varepsilon}'_{T+1,T+h} = \left(\boldsymbol{\varepsilon}'_{T+1} \dots \boldsymbol{\varepsilon}'_{T+h}\right)$ are the shocks over the forecasting horizon; we will refer to these as the future shocks. The vector $\mathbf{b}_{T+1,T+h}$ is predetermined, representing the dynamic forecast in the absence of further shocks. It depends only on the history of observables and the reduced-form parameters of the model. The vector $\mathbf{M}'\boldsymbol{\varepsilon}_{T+1,T+h}$ is stochastic as it involves the future shocks, and the matrix $\mathbf{M}$ depends on the structural parameters. Definitions of $\mathbf{b}_{T+1,T+h}$ and $\mathbf{M}'$ are given in the online Appendix. Given Equation (3), the unconditional forecast is distributed

$$\mathbf{y}_{T+1,T+h} \sim \mathcal{N}\left(\mathbf{b}_{T+1,T+h}, \mathbf{M}'\mathbf{M}\right). \tag{4}$$

It is easy to show that while $\mathbf{M}$ depends on the structural parameters, i.e., on the choice of $\mathbf{Q}$, $\mathbf{M}'\mathbf{M}$ only depends on the reduced-form parameters. Thus, one only needs the history of observables and the reduced-form parameters to characterize the distribution of the unconditional forecast. It is also easy to see that, irrespective of the choice of $\mathbf{Q}$, the future shocks are distributed according to their unconditional distribution $\boldsymbol{\varepsilon}_{T+1,T+h} \sim \mathcal{N}\left(\mathbf{0}_{nh\times 1}, \mathbf{I}_{nh}\right)$.

2.2. Conditional forecasts and structural scenarios: General framework

In general, linear restrictions on the path of future observables can be written as

$$\mathbf{C}\tilde{\mathbf{y}}_{T+1,T+h} \sim \mathcal{N}\left(\mathbf{f}_{T+1,T+h}, \boldsymbol{\Omega}_f\right), \tag{5}$$

where $\tilde{\mathbf{y}}_{T+1,T+h}$ is the restricted forecast of the observables, and $\mathbf{C}$ is a $k \times nh$ (full rank) pre-specified matrix, where k is the number of restrictions. The $k \times 1$ vector $\mathbf{f}_{T+1,T+h}$ and the $k \times k$ matrix $\mathbf{\Omega}_f$ are the mean and covariance matrix restrictions.[4] Equation (5) implies restrictions on the distribution of future shocks; combine Equations (3) and (5) to obtain

$$\mathbf{C}\tilde{\mathbf{y}}_{T+1,T+h} = \mathbf{C}\mathbf{b}_{T+1,T+h} + \mathbf{D}\tilde{\boldsymbol{\varepsilon}}_{T+1,T+h} \sim \mathcal{N}\left(\mathbf{f}_{T+1,T+h}, \mathbf{\Omega}_f\right), \tag{6}$$

where $\mathbf{D} = \mathbf{C}\mathbf{M}'$. Let $\tilde{\boldsymbol{\varepsilon}}_{T+1,T+h}$, the restricted future shocks, be distributed

$$\tilde{\boldsymbol{\varepsilon}}_{T+1,T+h} \sim \mathcal{N}\left(\boldsymbol{\mu}_\varepsilon, \mathbf{\Sigma}_\varepsilon\right) \tag{7}$$

so that, defining $\mathbf{\Sigma}_\varepsilon = \mathbf{I}_{nh} + \mathbf{\Psi}_\varepsilon$, $\boldsymbol{\mu}_\varepsilon$ and $\mathbf{\Psi}_\varepsilon$ are the deviation of the mean and covariance matrix of $\tilde{\boldsymbol{\varepsilon}}_{T+1,T+h}$ from their unconditional counterparts. Equation (6) implies the following restrictions on $\boldsymbol{\mu}_\varepsilon$ and $\mathbf{\Psi}_\varepsilon$

$$\mathbf{C}\mathbf{b}_{T+1,T+h} + \mathbf{D}\boldsymbol{\mu}_\varepsilon = \mathbf{f}_{T+1,T+h} \tag{8}$$

$$\mathbf{D}\left(\mathbf{I}_{nh} + \mathbf{\Psi}_\varepsilon\right)\mathbf{D}' = \mathbf{\Omega}_f. \tag{9}$$

Depending on k, the number of restrictions, and nh, the length of $\tilde{\mathbf{y}}_{T+1,T+h}$, the systems of Equations (8) and (9) may have multiple solutions ($k < nh$), one solution ($k = nh$), or no solutions ($k > nh$).[5] Following Penrose (1955, 1956) we will choose the following general expression for $\boldsymbol{\mu}_\varepsilon$ and $\mathbf{\Psi}_\varepsilon$,

$$\boldsymbol{\mu}_\varepsilon = \mathbf{D}^\star\left(\mathbf{f}_{T+1,T+h} - \mathbf{C}\mathbf{b}_{T+1,T+h}\right) \tag{10}$$

$$\mathbf{\Psi}_\varepsilon = \mathbf{D}^\star\mathbf{\Omega}_f\mathbf{D}^{\star\prime} - \mathbf{D}^\star\mathbf{D}\mathbf{D}'\mathbf{D}^{\star\prime}, \tag{11}$$

[4] This formulation accommodates, in the special case where $\mathbf{\Omega}_f = \mathbf{0}_{k\times k}$, the classic "hard" conditional forecasting exercise, as defined in Waggoner and Zha (1999), as well as its generalization to density conditional forecasts, as defined in Andersson et al. (2010).

[5] The presence of multiple solutions in Waggoner and Zha (1999) is highlighted by Jarociński (2010).

where $\mathbf{D}^{\star}$ is the Moore-Penrose inverse of $\mathbf{D}$. Because $\mathbf{D}$ is full rank, the Moore-Penrose inverse always exists and is unique. From the definition of $\boldsymbol{\Psi}_{\varepsilon}$, Equation (11) implies

$$\boldsymbol{\Sigma}_{\varepsilon} = \mathbf{D}^{\star}\boldsymbol{\Omega}_{f}\mathbf{D}^{\star\prime} + \left(\mathbf{I}_{nh} - \mathbf{D}^{\star}\mathbf{D}\mathbf{D}'\mathbf{D}^{\star\prime}\right), \tag{12}$$

a semidefinite positive matrix. Equations (10) and (12) make clear that linear restrictions on the path of future observables restrict the distribution of the future shocks.[6] Importantly, it can be shown that, given the reduced-form parameters, the choice of $\mathbf{Q}$ affects $\boldsymbol{\mu}_{\varepsilon}$ and $\boldsymbol{\Sigma}_{\varepsilon}$.

When $k \leq nh$, the system of Equations (8) and (9) is consistent, and Equations (10) and (11) characterize the solution that minimizes the Frobenius norm of $\boldsymbol{\mu}_{\varepsilon}$ and $\boldsymbol{\Psi}_{\varepsilon}$. Therefore, the Penrose solution is the one that envisages the smallest deviation of the mean and covariance matrix of $\tilde{\boldsymbol{\varepsilon}}_{T+1,T+h}$ from the mean and covariance matrix of $\boldsymbol{\varepsilon}_{T+1,T+h}$. Clearly, if $k = nh$ we have that $\mathbf{D}^{\star} = \mathbf{D}^{-1}$ and there is a unique solution, whereas when $k > nh$, the system is inconsistent, i.e., there is no solution to the system defined by Equations (8) and (9) that satisfies all the restrictions simultaneously. In this case Equations (10) and (11) are the best approximated solutions (see Penrose, 1956), meaning that they minimize

$$\left\|\mathbf{C}\mathbf{b}_{T+1,T+h} + \mathbf{D}\boldsymbol{\mu}_{\varepsilon} - \mathbf{f}_{T+1,T+h}\right\| \text{ and } \left\|\mathbf{D}\left(\mathbf{I}_{nh} + \boldsymbol{\Psi}_{\varepsilon}\right)\mathbf{D}' - \boldsymbol{\Omega}_{f}\right\|,$$

respectively, where we are using the Frobenius norm again.[7]

Let the restricted forecast of the observables be distributed $\tilde{\mathbf{y}}_{T+1,T+h} \sim \mathcal{N}\left(\boldsymbol{\mu}_{y}, \boldsymbol{\Sigma}_{y}\right)$. Then, given that $\tilde{\mathbf{y}}_{T+1,T+h} = \mathbf{b}_{T+1,T+h} + \mathbf{M}'\tilde{\boldsymbol{\varepsilon}}_{T+1,T+h}$, Equations (10) and (12) imply the following

[6] The reader should notice that the solutions in Equations (10) and (11) do not depend on our normality assumption. In particular, if one had assumed Student-t distributions for the shocks, the restrictions will still imply Equations (8) and (9) and, hence, the solutions in Equations (10) and (11) will still be valid.

[7] Specialized formulae for the three cases are given in the online Appendix. For inconsistent systems the Penrose solution gives the same weight to all the constraints. If a researcher wanted to give different weight to the different restrictions, the solution can be easily amended (see Ben-israel and Greville, 2001, p.117).

expressions for $\boldsymbol{\mu}_y$ and $\boldsymbol{\Sigma}_y$

$$\boldsymbol{\mu}_y = \mathbf{b}_{T+1,T+h} + \mathbf{M}'\mathbf{D}^{\star}(\mathbf{f}_{T+1,T+h} - \mathbf{C}\mathbf{b}_{T+1,T+h}) \tag{13}$$

$$\boldsymbol{\Sigma}_y = \mathbf{M}'\mathbf{M} + \mathbf{M}'\mathbf{D}^{\star}(\boldsymbol{\Omega}_f - \mathbf{D}\mathbf{D}')\mathbf{D}^{\star\prime}\mathbf{M}. \tag{14}$$

It can be shown that the distribution of the restricted forecast of the observables, given by Equations (13) and (14), only depends on the history of observables and the reduced-form parameters, unless $\mathbf{C}$ depends on $\mathbf{Q}$ (see Waggoner and Zha, 1999, Proposition 1).

Equations (13) and (14) highlight that the unconditional forecast is updated in order to incorporate the information contained in the restrictions.[8] The update is affected by how $\mathbf{f}_{T+1,T+h}$ and $\boldsymbol{\Omega}_f$ deviate from their unconditional mean and variance, $\mathbf{C}\mathbf{b}_{T+1,T+h}$ and $\mathbf{D}\mathbf{D}'$ respectively. Importantly, Equation (14) also highlights that choosing $\boldsymbol{\Omega}_f = \mathbf{D}\mathbf{D}'$ leads to $\boldsymbol{\Sigma}_y = \mathbf{M}'\mathbf{M}$, i.e., the variance of the unconditional forecast is preserved. This is a natural choice when one wants to restrict the mean of the forecast without affecting its variance. As we will see in the empirical examples, preserving the variance is of great importance if one is going to use the model for density forecasting and risk assessment.

Relation to entropic tilting. Finally, we provide a proof, in the Gaussian case, of the equivalence between the above framework and the method of entropic forecast tilting, popularized in the macroeconomic literature by Robertson et al. (2005) (see also Giacomini and Ragusa, 2014). The equivalence unifies two approaches that have evolved separately and at times have been regarded as alternatives. Entropic forecast tilting finds the forecast distribution subject to some linear restrictions that minimizes the relative entropy with the unconditional forecast, as measured by the Kullback-Leibler (KL) divergence. Define the KL divergence from Q to P as $D_{\text{KL}}(P\|Q) = \int_X p \log(\frac{p}{q})\, \mathrm{d}\mu$, where P and Q are probability distributions over a set X and μ is any measure on X for which $p = \frac{\mathrm{d}P}{\mathrm{d}\mu}$ and $q = \frac{\mathrm{d}Q}{\mathrm{d}\mu}$ exist (meaning that p and q are absolutely continuous with respect to μ). We formally establish the following:

[8] Banbura et al. (2015) propose an alternative formulation of Equations (13) and (14) that uses the Kalman filter and smoother, which we extend to the structural scenario case in the online Appendix. The formulas in Equations (13) and (14) will be computationally more efficient than the Kalman filter implementation for the majority of practical cases, except when the forecast horizon, h, is very large.

Proposition 1. *Denote with $\mathcal{N}_{UF}$ the distribution of the unconditional forecast represented by Equation (4). Then $\boldsymbol{\mu}_y$ and $\boldsymbol{\Sigma}_y$, given by Equations (13) and (14), are the solution to the following relative entropy problem*

$$\min_{\boldsymbol{\mu},\boldsymbol{\Sigma}} D_{KL}\left(\mathcal{N}\left(\boldsymbol{\mu},\boldsymbol{\Sigma}\right)||\mathcal{N}_{UF}\right)$$

subject to $\mathbf{C}\boldsymbol{\mu} = \mathbf{f}_{T+1,T+h}$ *and* $\mathbf{C}\boldsymbol{\Sigma}\mathbf{C}' = \boldsymbol{\Omega}_f$.

Proof 1. *See the online Appendix.*

The following corollary establishes the connection between the conditional forecasting solution and the entropic tilting problem that matches only the mean of the target distribution.

Corollary 1. *The conditional forecasting solution when $\boldsymbol{\Omega}_f = \mathbf{D}\mathbf{D}'$ is the solution to the relative entropy problem subject to the constraint $\mathbf{C}\boldsymbol{\mu} = \mathbf{f}_{T+1,T+h}$.*

In other words, conditional forecasting is equivalent to entropic tilting of the mean only when the uncertainty around the conditioning path is set to the unconditional variance, as discussed above. The close connection between conditional forecasting and entropic tilting motivates the use of the KL divergence to assess the plausibility of structural scenarios, an issue to which we return in Section 3.

2.3. Special Case 1: Conditional-on-observables forecasting

We now analyze three special cases of the framework presented above. The first is the classic conditional forecasting exercise, as first introduced by Doan et al. (1986) and extended by Waggoner and Zha (1999) and Andersson et al. (2010). It restricts the path of a subset of future observables and computes the forecast of the remaining observables.

Let $\overline{\mathbf{C}}$ be a $k_o \times nh$ (full rank) pre-specified selection matrix formed by ones and zeros, with k_o denoting the number of restrictions. Conditional-on-observables restrictions are written as $\overline{\mathbf{C}}\tilde{\mathbf{y}}_{T+1,T+h} = \overline{\mathbf{f}}_{T+1,T+h}$ (see Waggoner and Zha, 1999) or more generally as density restrictions $\overline{\mathbf{C}}\tilde{\mathbf{y}}_{T+1,T+h} \sim \mathcal{N}\left(\overline{\mathbf{f}}_{T+1,T+h}, \overline{\boldsymbol{\Omega}}_f\right)$ (see Andersson et al., 2010). These types of restrictions are trivially expressed in terms of Equation (5), by making $k = k_0$, $\mathbf{C} = \overline{\mathbf{C}}$, $\mathbf{f}_{T+1,T+h} = \overline{\mathbf{f}}_{T+1,T+h}$

and $\mathbf{\Omega}_f = \overline{\mathbf{\Omega}}_f$ and, hence, the expressions for $\boldsymbol{\mu}_\varepsilon$, $\mathbf{\Sigma}_\varepsilon$, $\boldsymbol{\mu}_y$, and $\mathbf{\Sigma}_y$ can be obtained using Equations (10), (12), (13), and (14) respectively.

Observe that in this case the selection matrix $\overline{\mathbf{C}}$ does not depend on the parameters of the model and that, hence, the distribution of $\tilde{\mathbf{y}}_{T+1,T+h}$ only depends on the history of observables and the reduced-form parameters. However, it is still the case that $\mathbf{Q}$ affects the distribution of the restricted future shocks. Given a particular $\mathbf{Q}$, one could back out such a distribution using Equations (10) and (12) and attach an economic interpretation to the conditional-on-observables forecast.

2.4. Special Case 2: Conditional-on-shocks forecasting

An alternative approach is to construct forecasts of observables restricting the path of a subset of shocks over the forecasting horizon.[9] Let $\mathbf{\Xi}$ be a $k_s \times nh$ (full rank) pre-specified selection matrix formed by ones and zeros, with k_s denoting the number of restrictions. Restrictions on the shocks can generally be written as $\mathbf{\Xi}\tilde{\boldsymbol{\varepsilon}}_{T+1,T+h} \sim \mathcal{N}\left(\mathbf{g}_{T+1,T+h}, \mathbf{\Omega}_g\right)$, where the $k_s \times 1$ vector $\mathbf{g}_{T+1,T+h}$ and the conformable matrix $\mathbf{\Omega}_g$ are the mean and covariance matrix restrictions to $\mathbf{\Xi}\tilde{\boldsymbol{\varepsilon}}_{T+1,T+h}$.[10] Provided that the VAR is invertible (see Fernández-Villaverde et al., 2007), the shocks can always be expressed as a function of observed variables, given the structural parameters of the model: $\tilde{\boldsymbol{\varepsilon}}_{t+1,t+h} = (\mathbf{M}')^{-1}\tilde{\mathbf{y}}_{t+1,t+h} - (\mathbf{M}')^{-1}\mathbf{b}_{t+1,t+h}$. Therefore, restrictions on future shocks can be written as linear restrictions on the path of future observables; specifically the restriction implies $\underline{\mathbf{C}}\tilde{\mathbf{y}}_{T+1,T+h} \sim \mathcal{N}\left(\underline{\mathbf{f}}_{T+1,T+h}, \underline{\mathbf{\Omega}}_f\right)$ where $\underline{\mathbf{C}} = \mathbf{\Xi}(\mathbf{M}')^{-1}$, $\underline{\mathbf{f}}_{T+1,T+h} = \underline{\mathbf{C}}\mathbf{b}_{T+1,T+h} + \mathbf{g}_{T+1,T+h}$ and $\underline{\mathbf{\Omega}}_f = \mathbf{\Omega}_g$. Thus, conditional-on-shocks forecasting can be expressed in terms of Equation (5) by making $k = k_s$, $\mathbf{C} = \underline{\mathbf{C}}$, $\mathbf{f}_{T+1,T+h} = \underline{\mathbf{f}}_{T+1,T+h}$ and $\mathbf{\Omega}_f = \underline{\mathbf{\Omega}}_f$ and, hence, the expressions for $\boldsymbol{\mu}_\varepsilon$, $\mathbf{\Sigma}_\varepsilon$, $\boldsymbol{\mu}_y$, and $\mathbf{\Sigma}_y$ can be obtained using Equations (10), (12), (13), and (14) respectively. The crucial point is that the matrix $\underline{\mathbf{C}}$ now depends on $\mathbf{M}$, which in turn depends on $\mathbf{Q}$. The intuition is clear; in order to impose restrictions upon their future path, we need to identify the shocks.[11]

[9] Baumeister and Kilian (2014) use an SVAR of the oil market to analyze the impact of a hypothetical oil supply shock. This practice is also used within DSGE models (see Del Negro and Schorfheide, 2013).

[10] Exact restrictions as in Baumeister and Kilian (2014) can be implemented fixing $\mathbf{\Omega}_g = \mathbf{0}_{k_s \times k_s}$.

[11] It is easy to show that, since the shocks are independent, the future shocks that are not part of the conditioning exercise retain their unconditional distribution.

2.5. Special Case 3: Structural scenario analysis

Conditional-on-shocks forecasting has the unappealing feature that, since the shocks are unobserved, it is difficult to elicit a priori restrictions on their future paths.[12] Here we show how the results of Sections 2.3 and 2.4 can be combined to approach this problem in a single step, which we call structural scenario analysis. A structural scenario is defined by combining restrictions on the path of future observables with a restriction that only a subset of the future shocks –the driving shocks– can deviate from their unconditional distribution over the forecasting horizon. Unlike in the conditional-on-observables case, the remaining future shocks –the non-driving shocks– are restricted to retain their unconditional distribution.

Let $\overline{\mathbf{C}}$ be a $k_o \times nh$ (full rank) pre-specified selection matrix formed by ones and zeros, with k_o denoting the number of restrictions on the observables. Let $\mathbf{\Xi}$ be a $k_s \times nh$ (full rank) pre-specified selection matrix formed by ones and zeros that selects the k_s non-driving shocks and whose distribution is going to be restricted to be the same as their unconditional one. As in Section 2.3, the restriction on the path of future observables is implemented by imposing $\overline{\mathbf{C}}\tilde{\mathbf{y}}_{T+1,T+h} \sim \mathcal{N}\left(\mathbf{f}_{T+1,T+h}, \mathbf{\Omega}_f\right)$., whereas as in Section 2.4, we impose that $\mathbf{\Xi}\tilde{\boldsymbol{\varepsilon}}_{T+1,T+h} \sim \mathcal{N}\left(\mathbf{0}_{k_s\times 1}, \mathbf{I}_{k_s}\right)$. The latter, in turn, implies $\underline{\mathbf{C}}\tilde{\mathbf{y}}_{T+1,T+h} \sim \mathcal{N}\left(\underline{\mathbf{C}}\mathbf{b}_{T+1,T+h}, \mathbf{I}_{k_s}\right)$, where $\underline{\mathbf{C}} = \mathbf{\Xi}(\mathbf{M}')^{-1}$. Thus, the restrictions embedded in a structural scenario can also be expressed in terms of linear restrictions on the path of future observables

$$\widehat{\mathbf{C}}\mathbf{y}_{T+1,T+h} \sim \mathcal{N}\left(\underbrace{\begin{bmatrix} \mathbf{f}_{T+1,T+h} \\ \underline{\mathbf{C}}\mathbf{b}_{T+1,T+h} \end{bmatrix}}_{\widehat{\mathbf{f}}_{T+1,T+h}}, \underbrace{\begin{bmatrix} \mathbf{\Omega}_f & \mathbf{0}_{k_o\times k_s} \\ \mathbf{0}_{k_s\times k_o} & \mathbf{I}_{k_s} \end{bmatrix}}_{\widehat{\mathbf{\Omega}}_f}\right), \tag{15}$$

with $\widehat{\mathbf{C}}' = [\overline{\mathbf{C}}', \underline{\mathbf{C}}']$. Observe that Equation (15) stacks the two sets of restrictions considered in previous subsections. The upper block states that a selection of variables must follow the path $\mathbf{f}_{T+1,T+h}$ in expectation; the second block states that the non-driving shocks must retain

[12] In practice, the papers that use that method calibrate the value of the shocks to generate the desired impact on a particular observable, or iterate between the shocks and the observables until achieving that result (see Baumeister and Kilian, 2014; Clark and McCracken, 2014).

their unconditional distribution. For the same reasons explained in Section 2.4, the structural scenario depends on the choice of identification scheme, given that the matrix $\widehat{\mathbf{C}}$ depends on $\mathbf{M}$, which in turn depends on $\mathbf{Q}$. The types of restrictions expressed in Equation (15) are again a special case of Equation (5), by making $k = k_o + k_s$, $\mathbf{C} = \widehat{\mathbf{C}}$, $\mathbf{f}_{T+1,T+h} = \widehat{\mathbf{f}}_{T+1,T+h}$ and $\mathbf{\Omega}_f = \widehat{\mathbf{\Omega}}_f$ and, hence, the expressions for $\boldsymbol{\mu}_\varepsilon$, $\mathbf{\Sigma}_\varepsilon$, $\boldsymbol{\mu}_y$, and $\mathbf{\Sigma}_y$ can be obtained using Equations (10), (12), (13), and (14) respectively.

The structural scenario analysis presented here is reminiscent of the counterfactual analysis often employed in the SVAR literature to report the impact of a particular structural shock assuming a specific path of a policy variable.[13] These counterfactuals correspond to particular structural scenarios where one variable and all shocks but one are restricted for the entire forecast horizon. Thus, the counterfactuals impose $k_o + k_s = nh$. The framework above is a generalization of such exercises, providing closed-form solutions for cases in which $k_o + k_s \neq nh$. Moreover, the previous literature has always considered the case with no uncertainty around the scenario, i.e., $\widehat{\mathbf{\Omega}}_f = \mathbf{0}_{(k_o+k_s)\times(k_o+k_s)}$, whereas the above framework generalizes to uncertain scenarios.[14] Finally, observe that in the specific case in which the restricted observable is a policy instrument and the only structural shock driving the scenario is the associated policy shock, our procedure is equivalent to the "modest policy interventions" analyzed by Leeper and Zha (2003). We expand on this point in the next subsection.

3. How plausible is the structural scenario?

When analyzing a structural scenario using SVARs, one should be careful that the implied shocks are not so "unusual" that the analysis risks falling into the criticism put forward by Lucas (1976). We previously highlighted how a structural scenario is associated with a distribution of the future driving shocks that deviates from its unconditional counterpart. When

[13] The objective of counterfactual analysis is to answer "what if" questions. For instance, Bernanke et al. (1997), Hamilton and Herrera (2004) and Kilian and Lewis (2011) analyze the impact of an oil price shock under the assumption that the Fed does not respond.

[14] The previous literature typically imposes $\widehat{\mathbf{\Omega}}_f = \mathbf{0}_k$ and, therefore, misrepresents the uncertainty in the scenarios. There are two types of uncertainty related to the fact that $\widehat{\mathbf{\Omega}}_f \neq \mathbf{0}_k$. By setting $\mathbf{\Omega}_f \neq \mathbf{0}_{k_o}$ we are considering uncertainty surrounding the restrictions on the observables. By setting the covariance matrix of $\mathbf{\Xi}\tilde{\boldsymbol{\varepsilon}}_{T+1,T+h}$ to $\mathbf{I}_{k_s}$ instead of to $\mathbf{0}_{k_s}$ we are considering uncertainty surrounding the non-driving future shocks.

this deviation is large, we deem the scenario implausible. Moreover, Proposition 1 established that Equations (13) and (14) are the solution that minimizes the relative entropy from the unconditional forecast to the structural scenario, or equivalently, from the unconditional distribution of the future shocks to their distribution under the scenario.[15] We therefore propose to use the KL divergence as a measure of how plausible a scenario is.

Specifically, denote with $\mathcal{N}_{SS}$ the distribution of the shocks compatible with the structural scenario and $\mathcal{N}_{UF}$ their unconditional distribution. Remember that $\mathcal{N}_{UF}$ is the standard normal distribution, so we have that the KL divergence from $\mathcal{N}_{UF}$ to $\mathcal{N}_{SS}$ is

$$D_{\text{KL}}(\mathcal{N}_{SS}\|\mathcal{N}_{UF}) = \frac{1}{2}\left(\text{tr}\left(\boldsymbol{\Sigma}_\varepsilon\right) + \boldsymbol{\mu}_\varepsilon'\boldsymbol{\mu}_\varepsilon - nh - \ln(\det \boldsymbol{\Sigma}_\varepsilon)\right) \tag{16}$$

where $\boldsymbol{\mu}_\varepsilon$ and $\boldsymbol{\Sigma}_\varepsilon$ are given by Equations (10) and (12).

While it is straightforward to compute $D_{\text{KL}}(\mathcal{N}_{SS}\|\mathcal{N}_{UF})$ using Equation (16), it is difficult to grasp whether any value for the KL divergence is large or small. In other words, the KL divergence can be easily used to rank scenarios but it is hard to understand how far away they are from the unconditional forecast. To ease the interpretation of the KL divergence, we adapt an idea of McCulloch (1989), who proposes "calibrating" the KL divergence from Q to P, two generic distributions, using the KL divergence between two easily interpretable distributions. In particular we suggest comparing $D_{\text{KL}}(\mathcal{N}_{SS}\|\mathcal{N}_{UF})$ with the divergence between two binomial distributions, one with probability q and the other with probability $1/2$.[16] Let $\mathcal{B}(m,p)$ denote the binomial distribution that runs m independent experiments, each of them with probability p of success. Here $m = nh$ represents the number of independent realizations of the shocks over the scenario. We suggest calibrating the KL divergence from $\mathcal{N}_{UF}$ to $\mathcal{N}_{SS}$ to a parameter q that would solve the following equation $D_{\text{KL}}(\mathcal{B}(nh,0.5)\|\mathcal{B}(nh,q)) = D_{\text{KL}}(\mathcal{N}_{SS}\|\mathcal{N}_{UF})$. The solution to the equation is $q = \frac{1}{2}\left(1+\sqrt{1-e^{-\frac{2z}{nh}}}\right)$ where $z = D_{\text{KL}}(\mathcal{N}_{SS}\|\mathcal{N}_{UF})$. In line

[15] The KL divergence, D_{KL}, is invariant to linear transformations, and hence it is equivalent to consider D_{KL} related to observables or shocks, since these are linearly related.

[16] McCulloch's (1989) original approach employed instead a comparison between two Bernouilli distributions. The drawback of that approach in our setting is that q is not invariant to the dimension of the scenario, nh. Specifically, for any $\boldsymbol{\mu}_\varepsilon \neq \mathbf{0}_{nh\times 1}$ and $\boldsymbol{\Sigma}_\varepsilon \neq \mathbf{I}_{nh}$, the KL divergence from $\mathcal{N}_{UF}$ to $\mathcal{N}_{SS}$ increases linearly with nh and $q \to 1$. The use of the binomial distribution retains the original interpretation but is scale invariant.

with McCulloch's (1989) original idea, any value for the KL divergence is translated into a comparison between the flip of a fair and a biased coin. For example, a value of $q = 0.501$ suggests that the distribution of the shocks under the scenario considered is not at all far from its unconditional counterpart, and the scenario considered is quite realistic. With a value of $q = 0.99$ the restrictions imply a substantial deviation of the future shocks from their unconditional distribution, suggesting that the scenario is quite unlikely.[17]

Our measure of plausibility is in the spirit of the concept of "modest intervention" used in Leeper and Zha (2003). Their measure reports how unusual is the path for policy shocks needed to achieve some conditional forecast, relative to the typical size of the shocks. If the scenario implies a sequence of shocks close to their typical size, the intervention is considered modest, and the results of the analysis are likely to be reliable. If instead the scenario involves an unlikely sequence of shocks, Leeper and Zha (2003) argue that economic agents may alter their behavior in meaningful ways and the analysis is deemed implausible.

Our metric compares the entire distribution rather than the path of the policy shocks, and generalizes to scenarios other than those involving a single policy instrument and a single policy shock.[18] Moreover, an appealing feature of a metric based on the KL divergence is that it provides an information-theoretical interpretation of the scenario, in terms of how much information is added by the restrictions. If a scenario involves meaningful distortions of the probability distribution of the underlying shocks, we should deem it as implausible and the metric calls for caution about the use of the SVAR for this particular task.

4. Algorithms for implementation

In this section we briefly sketch a Bayesian algorithm for implementation of the general framework above, which includes conditional-on-observables, conditional-on-shocks and

[17] In order to get a sense of how large is a specific calibration value, q, it might be useful to draw a parallel with a standard IRF, in which only one shock is active, equal to 1 s.d. for one period, and zero otherwise. With, say, $n = 8$ and $h = 12$, and assuming $\mathbf{\Sigma}_\varepsilon = \mathbf{I}_{nh}$, we have $q = 0.55$, a very reasonable scenario. A large, but not unreasonable, 2 s.d. shock, leads instead to $q = 0.6$. Moreover, a sequence of 1 s.d. shocks over 12 consecutive quarters leads to a $q = 0.67$, roughly equivalent to the q obtained from a single 3.5 standard deviation shock. For an extremely large shock of 10 s.d. we get a $q = 0.9$.

[18] That Leeper and Zha (2003) compare the path of the policy shocks is a consequence of the fact that they set all other shocks to zero.

structural scenarios as special cases. In the interest of space, we describe the main features here, relegating the technical details to the online Appendix. It should be clear that our objective is to draw from the following joint posterior $p(\tilde{\mathbf{y}}_{T+1,T+h}, \mathbf{A}_0, \mathbf{A}_+ | \mathbf{y}^T, \mathbf{IR}, \mathbf{SR})$ where **IR** are the structural identification restrictions, and **SR** are the restrictions on the path of future variables, shocks, or both. As explained by Waggoner and Zha (1999), drawing from this posterior distribution is a challenging task. It is tempting, in a first step, to draw the structural parameters from their distribution conditional on $\mathbf{y}^T$ and on **IR**, and in a second step to draw $\tilde{\mathbf{y}}_{T+1,T+h}$ conditional on $\mathbf{y}^T$, on **SR** and on the structural parameters using Equations (13) and (14). However, this procedure ignores the restrictions in Equation (15) when drawing the structural parameters and, hence, would not lead to a draw from the desired joint posterior. Instead, to draw from the joint distribution of interest, a Gibbs sampler procedure must be constructed that iterates between draws from the described conditional distributions of the structural parameters and $\tilde{\mathbf{y}}_{T+1,T+h}$ in the following way

Algorithm 1. *Initialize* $\mathbf{y}^{T+h,(0)} = [\mathbf{y}^T, \tilde{\mathbf{y}}^{(0)}_{T+1,T+h}]$*, e.g., at the mean unconditional forecast*

1. *Conditioning on* $\mathbf{y}^{T+h,(i-1)} = [\mathbf{y}^T, \tilde{\mathbf{y}}^{(i-1)}_{T+1,T+h}]$*, use any valid algorithm to produce one draw from the structural parameters,* $(\mathbf{A}_0^{(i)}, \mathbf{A}_+^{(i)})$ *satisfying* **IR**.

2. *Conditioning on* $(\mathbf{A}_0^{(i)}, \mathbf{A}_+^{(i)})$ *and* $\mathbf{y}^T$*, draw* $\tilde{\mathbf{y}}^{(i)}_{T+1,T+h} \sim \mathcal{N}(\boldsymbol{\mu}_y, \boldsymbol{\Sigma}_y)$ *satisfying* **SR** *using Equations* (13) *and* (14).

3. *Return to Step 1 until the convergence and a sufficient number of draws have been obtained.*

The algorithm above extends the one developed by Waggoner and Zha (1999), with the key difference that we now need to identify the SVAR. For example, in the case of the traditional recursive zero restrictions, Step 1 could draw the reduced-form coefficients and obtain $\mathbf{A}_0$ from the Cholesky decomposition of $\boldsymbol{\Sigma}$. In the case of sign restrictions, Step 1 may be any of the algorithms in Rubio-Ramirez et al. (2010) or Baumeister and Hamilton (2015). In the cases of narrative sign restrictions (NSR) as in Antolin-Diaz and Rubio-Ramirez (2018) and traditional sign and zero restrictions as in Arias et al. (2018), an importance-sampling step is required to draw from the posterior of $(\mathbf{A}_0, \mathbf{A}_+)$, which needs to be done at the end of Step 1.

4.1. The importance of using all available identification restrictions

When we use restrictions that set identify the model, the results of the structural scenario analysis will be valid across a set of structural models that satisfy the restrictions. This attractive feature will come at the cost of possibly wide confidence bands around the forecast. Most importantly, there is the risk of including many models with implausible implications for elasticities, structural parameters, shock realizations and historical decompositions.[19] For instance, in the applications with monetary policy shocks that we present below, we find that a strategy based exclusively on traditional sign restrictions on IRFs often leads to implausibly large elasticities of observable variables to a monetary policy shock, mirroring the results of Kilian and Murphy (2012).[20] In the examples, we will propose using a combination of traditional sign and zero restrictions at various horizons, as in Arias et al. (2016), bounds on elasticities, and the recently proposed NSR of Antolin-Diaz and Rubio-Ramirez (2018) to narrow down the set of admissible structural parameters and obtain meaningful scenarios.

5. Application 1: Monetary policy and the power of forward guidance

Our first application concerns the effects of communicating future monetary policy shocks. We provide new empirical estimates of the effects of anticipated monetary policy shocks, compare alternative structural scenarios to the ones obtained with DSGE models, and evaluate the potential of "makeup" strategies for inflation targeting.

Setup and identification. We work with the data set of Smets and Wouters (2007) (SW07), which contains seven key US macroeconomic time series at quarterly frequency: real output, consumption, investment, wages and hours worked per capita, inflation and the fed funds rate. Their data set starts in Q1 1966; we update it through Q4 2019. In order to identify anticipated monetary policy shocks, we will need a variable capturing changes in expectations of future interest rates. We follow an approach similar to that of Campbell et al. (2012) and

[19] This point has been forcefully argued by Kilian and Murphy (2012), Ahmadi and Uhlig (2015), Arias et al. (2016), Ludvigson et al. (2017) and Antolin-Diaz and Rubio-Ramirez (2018).

[20] These elasticities are much larger than the upper bound reported in Ramey's (2016) literature review. These translate into explosive forecasts of the variables even for modest deviations of the conditioned variable from its unconditional path.

use expectations from the Survey of Professional Forecasters (SPF), available for the first four quarters. Specifically, define as $E_t^{SPF}[i_{t+h}]$ the mean SPF expectation at time t of the short-term interest rate h quarters ahead. Then, the variable $S_{t,h} \equiv E_t^{SPF}[i_{t+h}] - E_{t-1}^{SPF}[i_{t+h}]$ for $h = 0, \ldots, 3$ measures the surprise to expectations, i.e., the change in the information set, from $t-1$ to t. Campbell et al. (2012) point out that surprises across horizons h exhibit a strong factor structure, so we take the average across horizons, $\bar{S}_t \equiv 1/4 \sum_{h=0}^{3} S_{t,h}$ as a measure of the surprises to the average expected interest rate for the next year.[21]. After appending $\bar{S}_t$ to the SW07 data set, we estimate an SVAR with $p = 4$ lags and an informative "Minnesota" prior (Doan et al., 1986; Sims, 1993) over the reduced-form parameters.[22] We identify two orthogonal monetary policy shocks by means of the following restrictions.

Identification Restriction 1: (*Unanticipated Monetary Policy Shocks*) The unanticipated monetary policy shock increases the fed funds rate on impact, decreases output, inflation, consumption, investment, hours and wages on impact, and has zero impact on the level of GDP in the long run, and a negative long-run impact on the price level. Moreover, it takes a positive value in Q4 1979 and is the single largest contributor to the unexpected movement in the fed funds rate in the same quarter. Finally, we restrict the maximum peak to trough response of output to a 100bp monetary policy shock to 10 pp.

The traditional sign and zero restrictions are satisfied by standard New Keynesian models such as Christiano et al. (1999) or SW07. The NSR follows Antolin-Diaz and Rubio-Ramirez (2018) and is motivated by the historical account of the Volcker tightening. The bound on the elasticity, set to twice the maximum amount reported by Ramey (2016), rules out models in which very small changes in the interest rate lead to massive declines in output.

[21] Using instead the first principal component yields essentially identical results. In the online Appendix, we show that inclusion of this variable among the observables in a DSGE model featuring anticipated shocks is enough to guarantee invertibility of the VAR representation, and thus the ability to recover the structural shocks from the reduced-form innovations (see Fernández-Villaverde et al., 2007)

[22] See the online Appendix for exact definitions of the variables and prior.

Identification Restriction 2: (*Anticipated Monetary Policy Shocks*) The anticipated monetary policy shock has the same sign effects on impact, zero long-run effects, and elasticity bounds as the unanticipated shock. Additionally, the effect on the fed funds rate is positive for the first three quarters, and it increases the surprise variable $\bar{S}_t$ on impact. Normalizing for the size of the impact on the fed funds rate, the increase in $\bar{S}_t$ is larger than the one caused by the unanticipated monetary policy shock. Finally, the anticipated monetary policy shock satisfies the NSR in Table 1.

Consistent with the evidence that policy surprises are correlated across horizons, we aim to identify a shock that affects the level of the interest rate over the next year, without imposing zero impact of the anticipated shock on the contemporaneous fed funds rate.[23] We derive the restrictions above from the ones implied by introducing an equivalently defined anticipated monetary policy shock within the SW07 DSGE model.[24] The NSR improve the identification by bringing in information about announcements made by the Federal Reserve in 2003 and 2011-2012.[25] Notice that we only impose the sign of the shocks because our reading of these events indicates that contemporaneous news about growth and inflation may have also influenced the announcements. We wish to stress that in general the surprise variable $\bar{S}_t$ is almost certainly affected in any quarter by many other shocks. Therefore, it would not constitute a valid instrument for anticipated monetary policy shocks. The advantage of the NSR approach in this context is that it allows us to incorporate the information that these shocks were affecting expectations during specific events.

For the results that follow it will be useful to also identify an "aggregate demand" shock, which generates positive co-movement between interest rates and business cycle variables.

[23] In our setting, an announcement of future actions with no contemporaneous change in the policy rate is best thought of as a combination of anticipated and unanticipated shocks

[24] Specifically, we introduce an anticipated monetary policy shock to the central bank reaction function in line with Campbell et al. (2012) and Campbell et al. (2017). The monetary policy 'news' shock obeys a simple factor structure for the first four quarters that embodies the idea that forward guidance is characterized by an announcement of a persistent deviation from the usual conduct of policy. Full details are available in the online Appendix.

[25] A motivation based on FOMC statements, minutes and transcripts is provided in the online Appendix.

Table 1: Summary of Forward Guidance FOMC Statement Language

	Economic Conditions		Language	NSR
	Growth	Inflation		
Q3 2003	"upside and downside risks [...] roughly equal"	"risk of inflation becoming undesirably low"	"policy accommodation can be maintained for a considerable period"	–
Q1 2004	"upside and downside risks [...] roughly equal"	"probability of an unwelcome fall[...] almost equal to that of a rise"	"patient in removing its policy accommodation"	+
Q2 2004	"upside and downside risks [...] roughly equal"	"risks to [...] price stability moved into balance"	"policy accommodation can be removed at a pace that is likely to be measured"	–
Q3 2011	"Considerably slower"	"picked up"	"low levels warranted at least through mid 2013"	–
Q1 2012	"expanding moderately"	"at or below mandate"	"low levels warranted at least through late 2014"	–
Q3 2012	"moderate pace"	"subdued"	"at least through mid 2015"	–

Identification Restriction 3: (*Aggregate Demand Shocks*) The aggregate demand shock has a positive effect on impact on output, consumption, investment, hours, wage, inflation and the fed funds rate, and zero long-run impact on the level of GDP.

Figure 2 displays the IRFs of key variables to the two policy shocks.[26] While the effects of the two shocks are quite similar, normalizing for the peak effect on the fed funds rate, the impact of the anticipated shock on output and inflation is somewhat larger and more persistent. The SVAR results do not display the "forward guidance puzzle" reported by Del Negro et al. (2012), and McKay et al. (2016), who find unrealistically large effects of anticipated shocks using DSGE models. We will expand on the reasons for this difference below.

Structural Scenario 1: Delphic vs. Odyssean forward guidance. The first of our exercises considers the two alternative questions we referred to in the introduction, the results of which we anticipated in Figure 1. In particular, we consider an exercise that might have been contemplated by the FOMC at the end of 2019. After three consecutive 25-basis-point cuts in previous meetings, the FOMC decided to keep the fed funds rate at a range of 1.5 to 1.75%. What if the FOMC had at this point announced a further reduction to 1%, after

[26] The IRFs of the remaining variables and shocks are displayed in the online Appendix.

Figure 2: Monetary Policy: Impulse Response Functions to anticipated vs unanticipated shocks

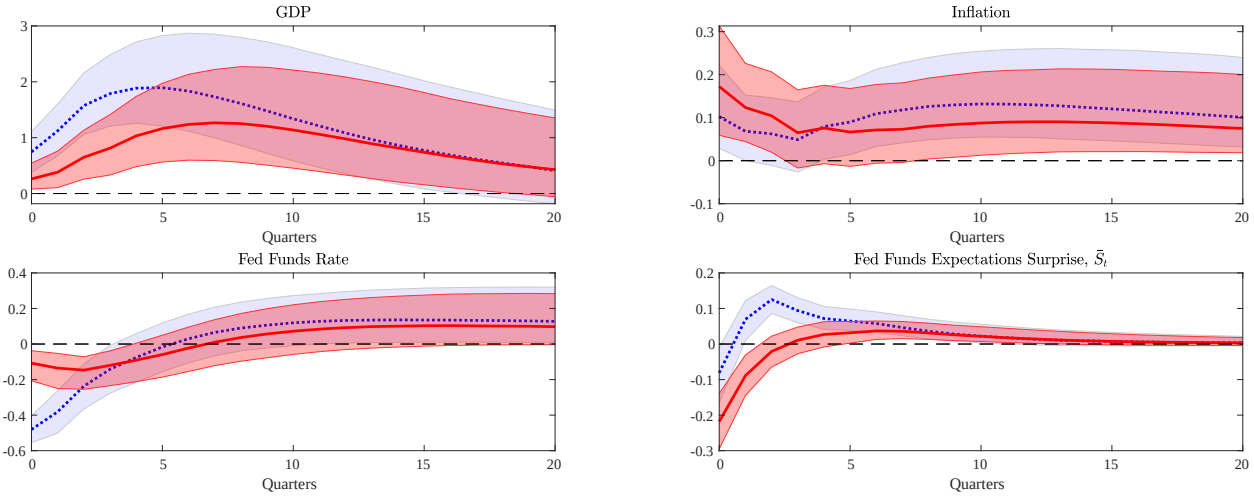

Note: The blue (lighter) shaded area represents the 68% (point-wise) HPD credible sets for the IRFs and the blue dotted lines are the median IRFs to an expansionary unanticipated monetary policy shock of one standard deviation. The red (darker) shaded areas and red solid lines display the equivalent quantities for the anticipated policy shock.

which the policy rate would be held constant for two years? In terms of our framework, this "dovish" scenario can be implemented as a restriction on the path of future observables using Equation (5). But whether we frame it in terms of a conditional-on-observables forecast or a structural scenario will yield completely different answers. Figure 3 illustrates this point by comparing the unconditional distribution of the three identified shocks in the first period of the forecast to the one resulting from several conditional-on-observables forecast and structural scenario exercises. As we can see in Panel (a), the conditional-on-observables exercise involves contractionary demand shocks generating, through the systematic reaction of the central bank, a positive co-movement between inflation and interest rates. As a result, output and inflation are weak. In this sense the exercise is at least partially Delphic. There are many occasions in which this setting may not be appropriate. For instance, in September 2012 the FOMC announced it expected to keep interest rates constant until mid-2015. The minutes of the meeting make clear that

"[the] new language was meant to clarify that the maintenance of a very low fed

Figure 3: Distribution of Shocks Underlying Different Exercises

(a) Conditional-on-Observables Forecast

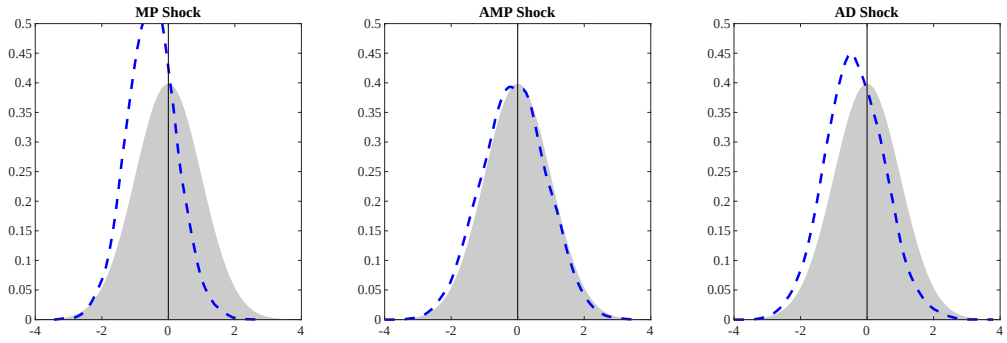

(b) An Odyssean Scenario

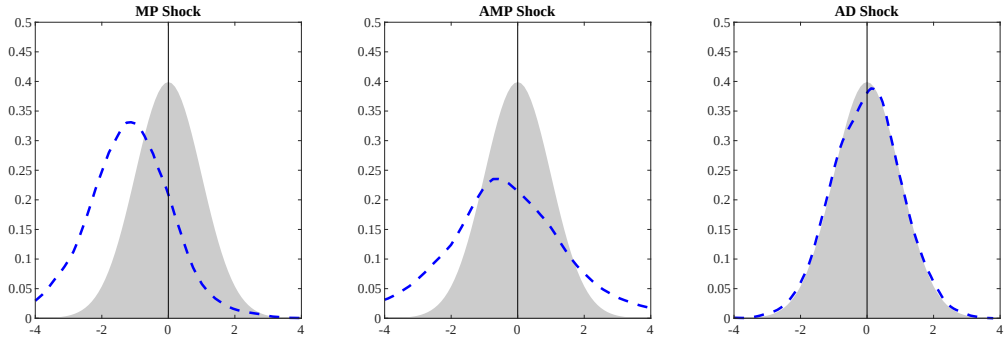

(c) A Delphic Scenario

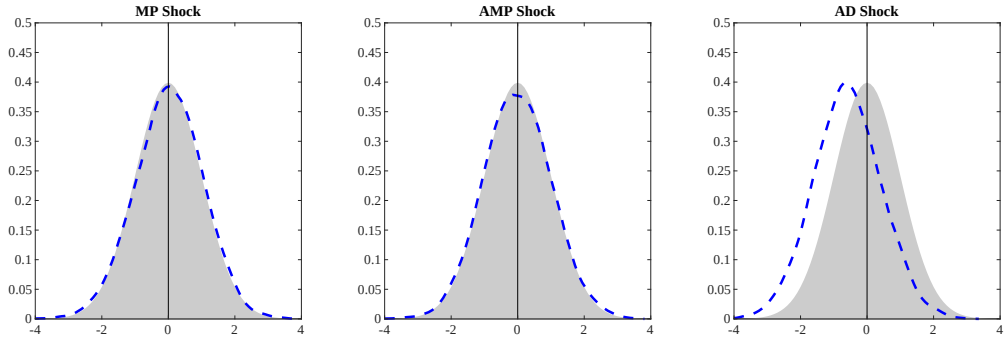

Note: For each panel, the gray shaded area represents the probability density function (p.d.f.) of a standard normal distribution, whereas the broken blue line is the p.d.f. of the different shocks at horizon $h = 1$ for the three variants of the "dovish" scenario. MP stands for unanticipated monetary policy shocks, AMP for anticipated monetary policy shocks, and AD for aggregate demand shocks.

> funds rate over that period did not reflect an expectation that the economy would remain weak, but rather reflected the Committee's intention to support a stronger economic recovery."[27]

We make this idea operational with a scenario in which all non-policy shocks are restricted to retain their unconditional distributions. By doing so, we avoid inferring anything about future demand and supply shocks from the lower path of interest rates. The results for the path of output and inflation are visible in Panel (b) of Figure 1, and for the associated shocks in Panel (b) of Figure 3. As we can see, there is a moderate boom in output, and inflation is slightly above the unconditional forecast. In the conditional-on-observables forecast, monetary policy shocks also played a role, reflecting the likelihood that at least some movements in interest rates represent exogenous policy actions. In that sense it mixes Delphic and Odyssean content. Can we construct a strictly Delphic scenario? The answer is yes. We create a Delphic variant of the scenario by restricting the two policy shocks to their unconditional distribution, leaving all other shocks unconstrained. The results can be seen in Figure 4 and the associated distribution of first period shocks is in Panel (c) of Figure 3. Output and inflation are now even lower than in the conditional-on-observables forecast, since we are shutting down the possibility that expansionary monetary policy shocks play any role.

The importance of quantifying uncertainty. In the scenarios above we fixed the path for the fed funds rate with zero uncertainty, i.e., $\mathbf{\Omega}_f = \mathbf{0}_{k_o \times k_o}$, as has been traditional in the literature. As pointed out by Andersson et al. (2010), this will lead to a reduction in the variance of the conditional forecast relative to the unconditional forecast. Instead, when one wants to preserve the uncertainty present in the unconditional forecast, the variance around the fed funds rate can be set to $\mathbf{\Omega}_f = \mathbf{D}\mathbf{D}'$, as proposed in Section 2.2. In the monetary policy example, this has an interesting economic interpretation. When fixing the policy instrument to a particular path, the scenario can be interpreted as an ironclad commitment: the monetary policy will follow that path no matter what the realizations of the other shocks

[27] Emphasis ours. See "Minutes of the Federal Open Market Committee" September 12-13, 2012 (Released October 4, 2012). https://www.federalreserve.gov/monetarypolicy/fomcminutes20120913.htm

Figure 4: A Delphic Scenario

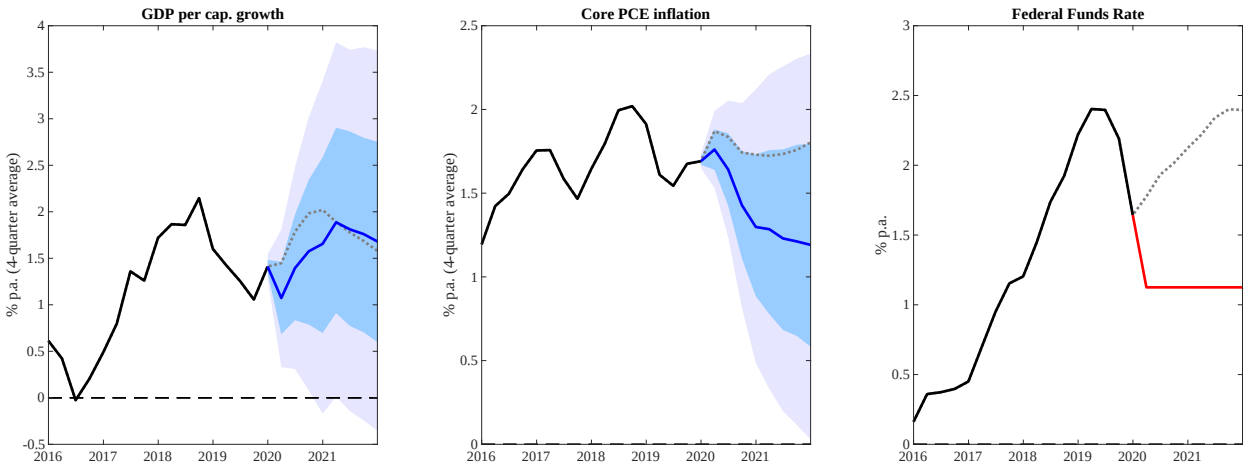

Note: For each panel, the solid black lines represent actual data, the solid red line is the conditioning assumption on the fed funds rate, the solid blue line is the median forecast for the remaining variables and periods, and the blue shaded areas denote the 40 and 68% pointwise credible sets around the forecasts. The dotted gray lines represent the median unconditional forecast.

are. If instead uncertainty about the fed funds path is recognized, the exercise acquires the flavor of a "data dependent" commitment.

If one further imposed that the "non-driving" shocks are equal to zero, i.e., $\mathbf{\Omega}_g = \mathbf{0}_{k_s}$, an even greater reduction in the variance of the forecast would result.[28] Intuitively, this assumption implies that the forecaster has knowledge that these shocks will not occur. As a result, the forecast density produced will greatly underestimate the underlying uncertainty about the future. Instead, by setting their distribution equal to the unconditional distribution, i.e., $\mathbf{\Omega}_g = \mathbf{I}_{k_s}$ we take the position of a forecaster who possesses no additional knowledge (over the unconditional forecast) about the "non-driving" shocks. Ignoring these two sources of uncertainty may make the results implausible and substantially less useful, particularly if one is interested in density forecasting. To illustrate this point, Table 2 reports, for the Odyssean variant of Scenario 1, the plausibility metric from Section 3 and a number of risk

[28] This is the implicit assumption of existing methodologies, including Leeper and Zha (2003), Baumeister and Kilian (2014), and Banbura et al. (2015). A figure displaying the consequences of the alternative treatments of uncertainty for the density forecasts for this exercise is available in the online Appendix.

Table 2: Risk Assessment of Structural Scenarios

	Calibrated KL	g at Risk	π at Risk	$p(\text{low } \pi, g)$	$p(g < 0)$
Unconditional	0.5	−1.2%	−0.7%	6.5%	26%
"Dovish - Odyssean" no uncertainty	1	1.8%	0.8%	0.1%	0.4%
"Dovish - Odyssean" with uncertainty	0.83	−2.0%	−1.5%	7.4%	26%
"Back to PL Target" with uncertainty	0.73	−2.4%	−1.3%	5.0%	36%

Note: The first column reports the posterior mode across draws of the calibrated K metric. The next columns report the growth-at-risk (g at Risk) and inflation-at-risk (π at Risk), defined respectively as the values associated with the 5% lower percentile of the predictive distribution for four-quarter GDP growth and inflation one year into the forecast horizon, as well as the probability of the average quarterly GDP growth rate and inflation rate in the first two years of the forecast both being below 1%($p(\text{low } \pi, g)$), and the probability of negative average quarterly growth in the third year of the forecast ($p(g < 0)$). The first row reports the results for the unconditional forecast, whereas the second row reports the case when the uncertainty of both the conditioning variable and the non-driving shocks is set to zero, i.e., $\mathbf{\Omega}_f = \mathbf{0}_{k_o \times k_o}$. The last three columns report results for scenarios where uncertainty about both is allowed, i.e., $\mathbf{\Omega}_f = \mathbf{D}\mathbf{D}'$.

metrics, such as "value at risk" for growth and inflation as well as the probability of observing a prolonged period of low growth and inflation with and without both sources of uncertainty. Notice that the posterior mode of the plausibility metric is equal to 1 (a large distortion of the distribution of shocks) and the risk metrics are considerably more sanguine than in the unconditional forecast when uncertainty is ignored. Instead, when uncertainty about both the conditioning variable and the non-driving shocks is recognized, the posterior mode of the plausibility metric is 0.83 (an unlikely but not completely implausible scenario) and the risk metrics are now an order of magnitude similar to the ones in the unconditional forecast.

A final observation relates to the fact that risks are generally deemed to be higher when considering structural scenarios than in the unconditional forecast. This reflects uncertainty stemming from the set identification of the structural parameters, which would be absent in exactly-identified models. In other words, with set identification the researcher is more uncertain about the causal impact of monetary policy shocks, whereas this uncertainty cancels out in the unconditional forecast, which is based only on correlations. This additional uncertainty needs to be taken into account when performing risk assessment exercises.

Structural Scenario 2: Evaluating forward guidance in the time of Greenspan. Our next exercise considers the difference between anticipated and unanticipated monetary shocks in more detail, and compares the results of our approach to the ones that would be obtained with the workhorse SW07 DSGE model. In particular, we examine the tightening cycle of 2004-2006, noteworthy for the fact that, in the prior year, the FOMC had been experimenting with forward guidance about the future path of interest rates. To simulate a real-time exercise we re-estimate our SVAR model using the vintage of data ending in Q4 2004, the original data set used by SW07.[29] For the DSGE model we fix the parameters to the posterior modal estimates reported by Smets and Wouters and introduce an anticipated monetary policy shock that is consistent with our identification strategy above and construct equivalent scenarios.[30] In Q3 2004, the FOMC initiated a tightening sequence that increased the fed funds rate by 25-basis-points each meeting until reaching 5.25%. What if instead the FOMC had decided to postpone that process by one year, keeping interest rates constant until then using policy shocks? How would the results differ depending on whether such an action was communicated in advance? Figure 5 compares the answers obtained using our SVAR and the DSGE model. Panel (a) implements this scenario only through unanticipated shocks. As we can see, the results of the SVAR are very similar to the point estimates from the DSGE model, though the effect on inflation appears slightly more front loaded. As we can see in Panel (b), things are more different in the case in which anticipated shocks are included. Using New Keynesian DSGE models, Del Negro et al. (2012) and McKay et al. (2016) find responses to anticipated policy shocks that are unrealistically large, and become even larger the longer the horizon of anticipation. The SVAR instead finds that the differences between anticipated and unanticipated policies are relatively minor.[31] Moreover, in the present exercise, since the horizon of anticipation spans only the first four quarters, we observe only a modest forward

[29] To identify the shocks, we only use the NSR up to Q2 2004.

[30] Since the SW07 model does not have a VAR representation, we cannot apply Equations (10)-(14) or algorithm 1 directly to compute the scenarios for the DSGE model. However, conditioning on the parameters, we can use the state-space representation of the model and the Kalman filter to compute the mean of the scenario. For this reason we only report point estimates for the SW07 model.

[31] The plausibility metric is 0.84 in the case where anticipation is allowed, and 0.8 when it is not, indicating that such exercises already imply a meaningful divergence from the unconditional distribution.

guidance puzzle with the DSGE, and the results are similar to the ones obtained with the SVAR. For scenarios that envision more drastic deviations from the unconditional forecast, or actions expanding into more distant horizons, however, the differences can become very large. Our empirical results provide a useful benchmark for any theoretical developments that aim to solve the forward guidance puzzle. The exercise also highlights one of the main practical motivations behind our paper: the scenarios conducted with DSGE models can be sensitive to the details of the model assumptions, whereas the ones constructed with SVARs will be consistent with a set of incompletely specified and not fully trusted models. When one is uncertain about which underlying structural model to use, our methodology will constitute a pragmatic alternative.

Structural Scenario 3: Returning to the price level target. Our final exercise considers a different thought experiment. Since the FOMC announced a 2% inflation target in January 2012, the average rate of inflation as measured by the PCE deflator excluding food and energy turned out to be 1.6%, resulting in a 4% cumulative shortfall relative to a hypothetical price level target by the end of 2019. The possibility of adopting a "makeup" strategy that would aim to revert the past shortfall has been widely discussed in policy circles in recent years.[32] What distribution of policy shocks would be required going forward to return to the 2% price level target, say, within 5 years? It is straightforward to implement this question using our structural scenario framework where only the two policy shocks are allowed to deviate from their unconditional distribution. Suppose that inflation is the i-th variable in the system and that one wanted to eliminate a cumulative shortfall of $\hat{\pi}$ over h periods. We can write this restriction as $1/h \sum_{s=1}^{h} y_{i,t+s} \sim \mathcal{N}\left(2+\hat{\pi}/h, \omega_f\right)$. The restriction states that the average inflation over the forecast horizon needs to be, in expectation, above the 2% target by $\hat{\pi}/h$ every period in order to eliminate the past accumulated inflation shortfall, with variance ω_f. Notice that our methodology is flexible enough that we do not need to input a specific path of inflation. Our method will yield the distribution of future policy shocks that is consistent with the smallest deviation from their unconditional distribution while satisfying

[32] See, e.g., Bernanke (2017). Those "makeup" strategies are also discussed within the Federal Reserve "Review of Monetary Policy Strategy, Tools, and Communication Practices;" see, e.g., Clarida (2019).

Figure 5: Comparison of SVAR and DSGE results

(a) Unanticipated Shocks Only

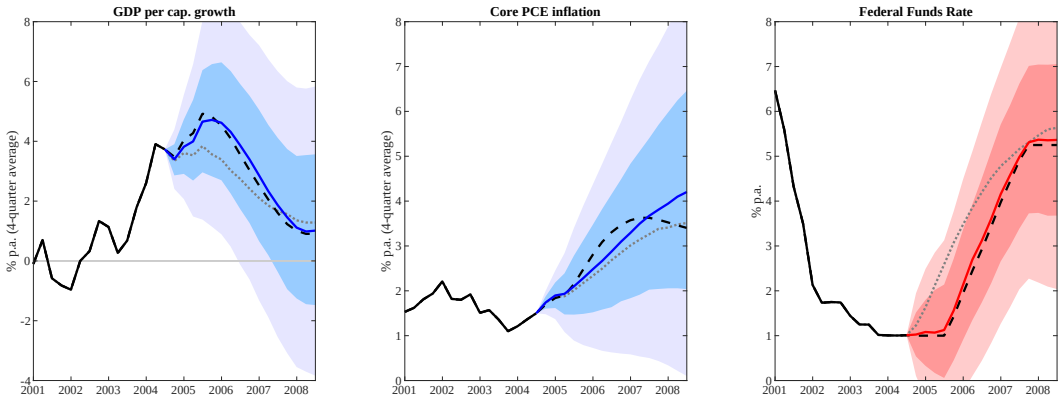

(b) Anticipated and Unanticipated Shocks

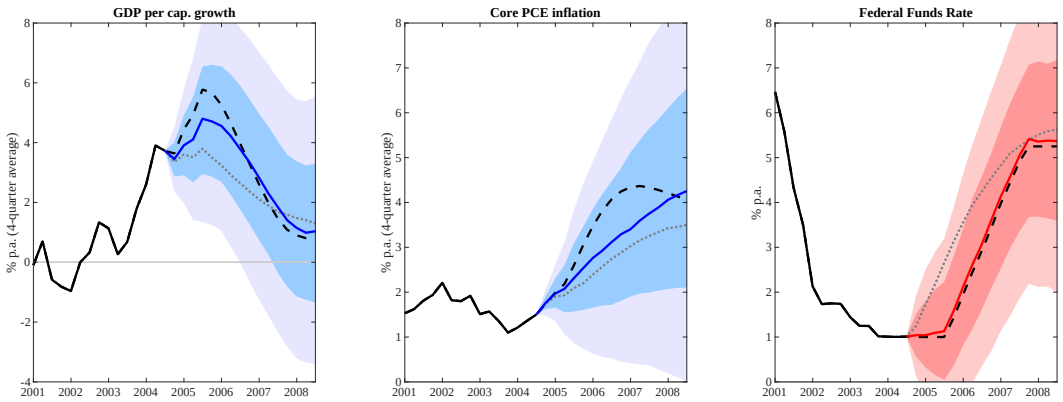

Note: For each panel, the solid black lines represent actual data. In the third column, the solid red line and associated shaded areas are the median, 40 and 68% pointwise credible sets around the conditioning assumptions for the fed funds rate. In the first and second panels, the solid blue line and associated shaded areas are the median, 40 and 68% pointwise credible sets around the forecasts of the other variables given by the SVAR. The dotted gray lines represent the median unconditional forecast from the SVAR. The broken black lines represent instead the mean results obtained using the DSGE model of SW07.

the restriction on $1/h \sum_{s=1}^{h} y_{i,t+s}$. Moreover, the interest rate distribution we obtain will be the one that "makes up" for the past inflation shortfall with the smallest deviation from the usual policy rule. The results can be seen in Figure 6. Inflation needs to increase gradually, peaking at 2.9% in mid-2023, before returning to the 2% target. In order to achieve the inflation, the fed funds rate is cut to about 1.3%, after which it starts increasing rapidly, until reaching 4% in 2023. This large tightening is required to cool down the boom in output

and inflation triggered by the new policy and bring inflation back to target. Using a car analogy, to catch up a car ahead of you that goes as fast as you, one needs to accelerate first and then decelerate. Decelerating the pace of price increases, requires higher interest rates to weaken aggregate demand. The results for output highlight a clear risk: while the price level successfully returns to the target, the subsequent tightening implies meaningful probabilities (around 46%) of negative annual growth in the fourth and fifth years of the sample. If the return to the price level target would be executed over a longer period, the implied slowdown would be less pronounced. Interestingly, this scenario produces larger effects on output and inflation than the apparently more aggressive interest rate cuts of Scenario 1. The reason is that Scenario 3 involves a greater share of anticipated monetary policy shocks relative to Scenario 1, which included a large unanticipated movement in the first period. As reported in Table 2, the plausibility metric for this scenario has a posterior mode of 0.73, suggesting it is a much more likely scenario than Scenario 1. Our results are consistent with policy communication being an effective way to achieve the central bank's objectives.[33]

6. Application 2: Stress-testing the response of asset prices to a recession

Our second application studies the impact of an economic recession on asset prices and bank profitability. We highlight that the potential impact of a recessionary episode can be very different depending on the shock driving it. Specifically, we will show how a recession caused by financial shocks, like the one of 2007-09, can have a more damaging impact on bank profitability (and other financial variables) than recessions driven by other types of shocks. We augment the macro data set from the previous application by adding the unemployment rate, as well as a number of financial variables: the 3-month Treasury bill to 10-year government bond yield spread, the quarterly change in the S&P 500 stock price index, the quarterly change in the S&P Case-Shiller House Price Index, the real price of oil, the BAA credit spread, the 3-month Treasury bill-eurodollar (TED) spread; and an indicator of profitability

[33]Two alternative scenarios are presented in the online Appendix. The first considers returning to the price level target over ten years instead of five. The second one features only unanticipated shocks. The results of the latter indicate that a more volatile interest rate path and a higher probability of recession would be required to achieve the same objective.

Figure 6: Returning to the Price Level Target

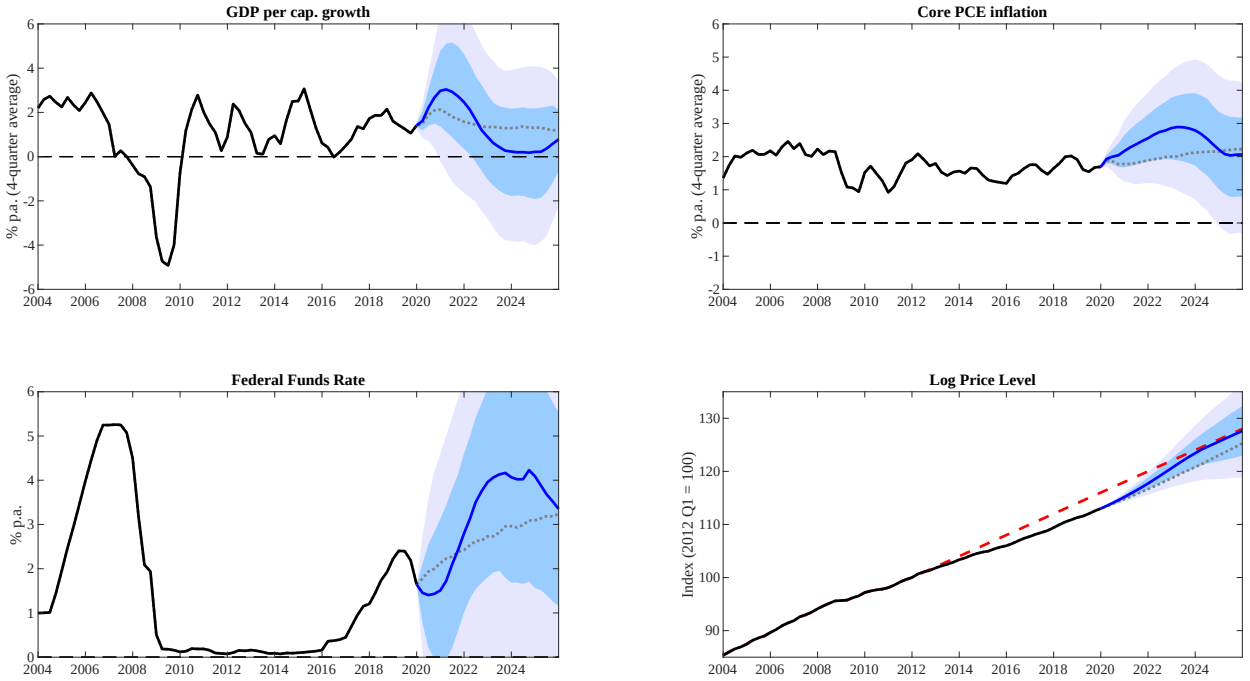

Note: For each panel, the solid black lines represent actual data. The solid blue lines and associated shaded areas are the median, 40 and 68% pointwise credible sets around the forecasts of the other variables given such that the constraint that the price level returns to the target within a five-year horizon is satisfied. The dotted gray lines represent the median unconditional forecast.

in the banking system as a whole, the return on equity (ROE) of FDIC-insured institutions.[34] Sample is from Q1 1966 to Q4 2019. We identify a single financial shock that resembles the one at work during the last recession through the following combination of traditional and narrative sign restrictions:

Identification Restriction 1: (*Financial Shock*) The financial shock has a negative impact on stock prices and bank profitability, and increases the BAA and TED spreads. Moreover, the financial shock for the observation corresponding to the fourth quarter of 2008 must be of positive value and be the overwhelming driver of the unexpected movement in the TED spread and credit spread.

[34] See the online Appendix for the definition of the variables and details of the model's specification.

The NSR are based on the account in Bernanke (2015), which highlights how the collapse of Lehman Brothers on September 13, 2008 caused "short-term lending markets to freeze and increase the panicky hoarding of cash" (p. 268), "fanned the flames of the financial panic" (p. 269), "directly touched off a run on money market funds" (p. 405), and "triggered a large increase in spreads" (p. 405). The construction of the scenarios borrows from the Federal Reserve's "2019 Supervisory Scenarios for Annual Stress Tests Required under the Dodd-Frank Act Stress Testing Rules and the Capital Plan Rule."[35] We take the path of GDP and the unemployment rate described in the Fed's "adverse scenario." It describes a medium-sized recession in which output falls for five consecutive quarters and then recovers gradually, whereas the unemployment rate increases until it reaches 7%. As before, we consider uncertainty around the path of the conditioning variables by setting $\mathbf{\Omega}_f = \mathbf{D}\mathbf{D}'$. Conditional on the restrictions for GDP and unemployment, we consider two distinct scenarios. In the first, the recession is caused by a financial shock. This is achieved by restricting all shocks except the financial one to their unconditional distribution. The second scenario is a recession not driven by the financial shock. To implement this, we restrict the financial shock to retain its unconditional distribution, and allow all the other shocks to deviate from theirs.

Figures 7 and 8 compare the results. The financial recession has a much more severe impact on stock prices, credit spreads and bank profitability. Moreover, the impact of the oil price is of opposite sign, indicating that non-financial recessions are often associated with increases in oil prices. As for the plausibility of the exercise, for the non-financial recession, the posterior mode of the calibrated KL measure is 0.68, whereas for the financial recession it is 0.86. We interpret this result as reflecting the fact that in the sample, financial recessions are a relatively infrequent event, and most postwar recessions appear to have involved a mixture of non-financial factors. Therefore, these results based on relatively infrequent events should be interpreted with caution. The results of this section highlight that if the scenario is driven by a financial shock, even mild contractions in economic activity can result in large damage to banks' balance-sheet variables.

[35] See https://www.federalreserve.gov/newsevents/pressreleases/files/bcreg20190213a1.pdf downloaded on January, 31, 2020.

Figure 7: A Stress Test Driven by a Financial Shock

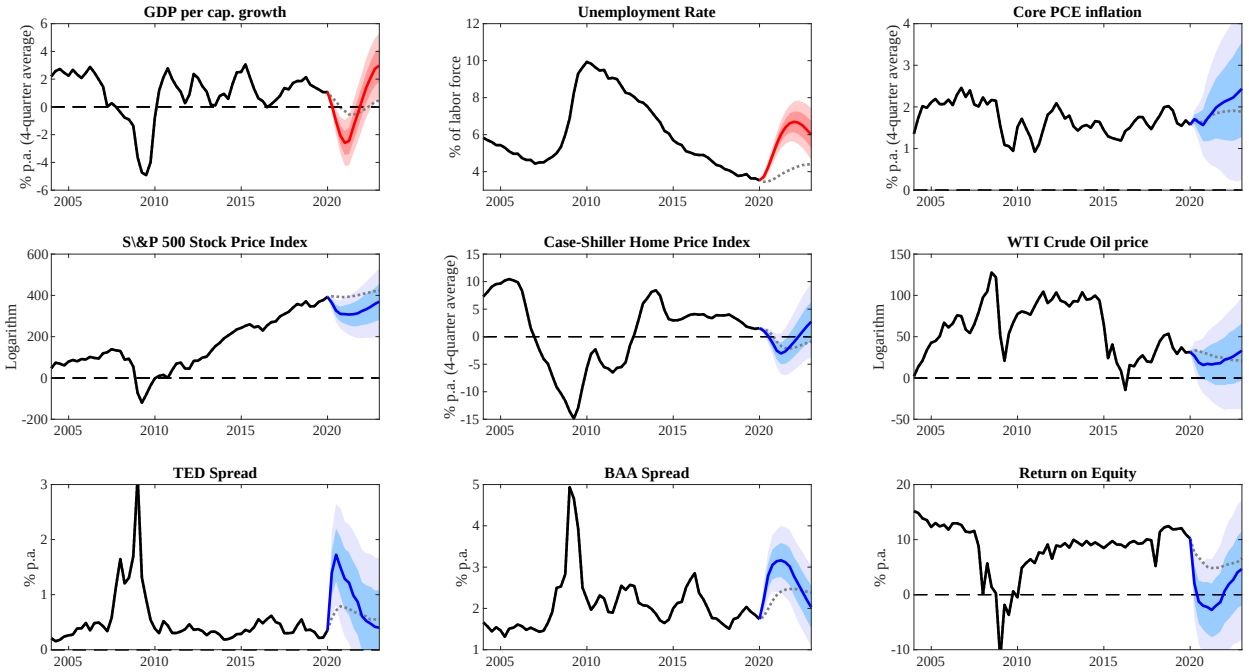

Note: For each panel, the solid black lines represent actual data. In the first two panels, the solid red line and associated shaded areas are the median, 40 and 68% pointwise credible sets around the conditioning assumptions for GDP growth and the unemployment rate. In the remaining panels, the solid blue line and associated shaded areas are the median, 40 and 68% pointwise credible sets around the forecasts of the other variables. The dotted gray lines represent the median unconditional forecast.

7. Conclusion

We have provided a general and flexible framework to construct structural scenarios based on minimalistically identified VARs, paying close attention to the uncertainty stemming from estimation, identification, and future shocks. While there will always be limits to the kinds of experiments that can be performed on linear Gaussian models, we hope that our methods facilitate a middle ground between relying on empirical models, which remain silent about the underlying causal mechanisms, and committing to the specific details of a particular structural model, which will give sharper, but surely misspecified answers.

References

Aastveit, K. A., A. Carriero, T. E. Clark, and M. Marcellino (2017): "Have standard VARs remained stable since the crisis?" *Journal of Applied Econometrics*, 32, 931–951.

Figure 8: A Stress Test Driven by Non-Financial Shocks

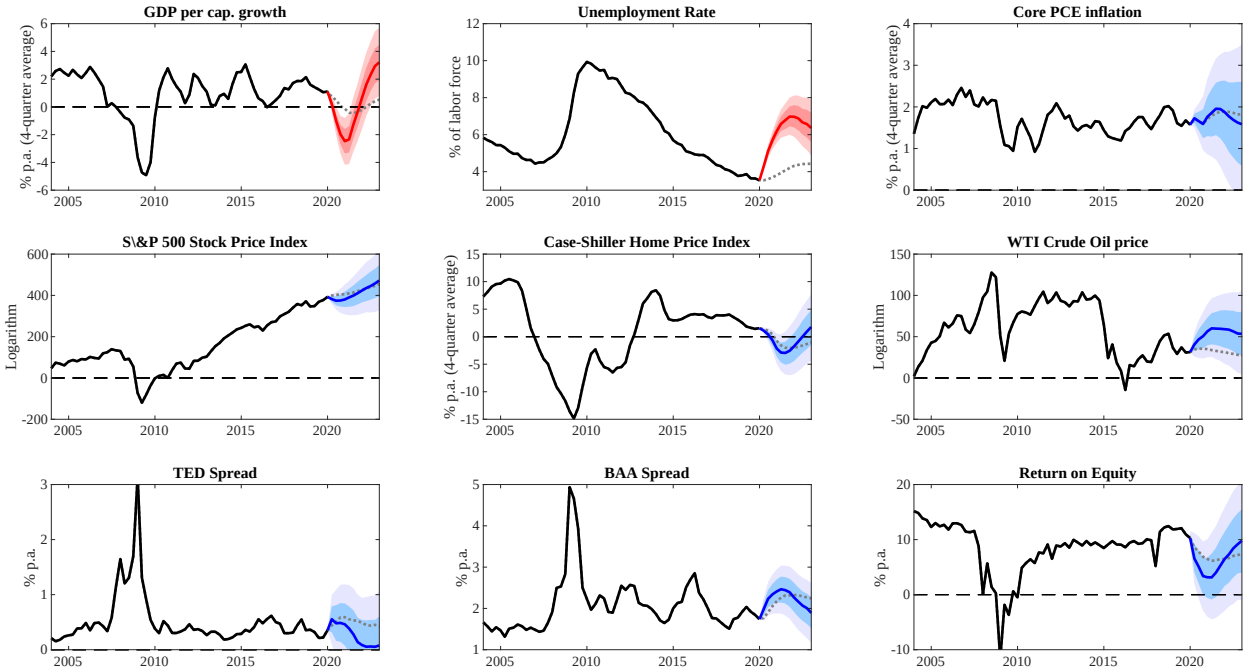

Note: For each panel, the solid black lines represent actual data. In the first two panels, the solid red line and associated shaded areas are the median, 40 and 68% pointwise credible sets around the conditioning assumptions for GDP growth and the unemployment rate. In the remaining panels, the solid blue line and associated shaded areas are the median, 40 and 68% pointwise credible sets around the forecasts of the other variables. The dotted gray lines represent the median unconditional forecast.

Ahmadi, P. A. and H. Uhlig (2015): "Sign Restrictions in Bayesian FaVARs with an Application to Monetary Policy Shocks," NBER Working Papers 21738, National Bureau of Economic Research, Inc.

Altavilla, C., D. Giannone, and M. Lenza (2016): "The Financial and Macroeconomic Effects of the OMT Announcements," *International Journal of Central Banking*, 12, 29–57.

Andersson, M. K., S. Palmqvist, and D. F. Waggoner (2010): "Density Conditional Forecasts in Dynamic Multivariate Models," Sveriges Riksbank Working Paper Series 247, Sveriges Riksbank.

Antolin-Diaz, J. and J. F. Rubio-Ramirez (2018): "Narrative Sign Restrictions for SVARs," *American Economic Review*, 108, 2802–29.

Arias, J. E., D. Caldara, and J. F. Rubio-Ramirez (2016): "The Systematic Component of Monetary Policy in SVARs: An Agnostic Identification Procedure," International Finance Discussion Papers 1131, Board of Governors of the Federal Reserve System (U.S.).

Arias, J. E., J. F. Rubio-Ramirez, and D. F. Waggoner (2018): "Inference Based on SVARs Identified with Sign and Zero Restrictions: Theory and Applications," *Econometrica*, 86, 685–720.

Banbura, M., D. Giannone, and M. Lenza (2015): "Conditional forecasts and scenario analysis with vector autoregressions for large cross-sections," *International Journal of Forecasting*, 31, 739–756.

Baumeister, C. and J. D. Hamilton (2015): "Sign restrictions, structural vector autoregressions, and useful prior information," *Econometrica*, 83, 1963–1999.

Baumeister, C. and L. Kilian (2014): "Real-Time Analysis of Oil Price Risks Using Forecast Scenarios," *IMF Economic Review*, 62, 119–145.

Ben-israel, A. and T. N. E. Greville (2001): *Generalized Inverses: Theory and Applications*, Springer-Verlag New York, second ed.

Bernanke, B. S. (2015): *The Courage to Act*, W.W. Norton & Company.

——— (2017): "Monetary Policy in a New Era," Tech. rep., Peterson Institute for International Economics.

Bernanke, B. S., M. Gertler, and M. Watson (1997): "Systematic Monetary Policy and the Effects of Oil Price Shocks," *Brookings Papers on Economic Activity*, 28, 91–157.

Campbell, J. R., C. L. Evans, J. D. Fisher, and A. Justiniano (2012): "Macroeconomic Effects of Federal Reserve Forward Guidance," *Brookings Papers on Economic Activity*, 43, 1–80.

Campbell, J. R., J. D. Fisher, A. Justiniano, and L. Melosi (2017): "Forward guidance and macroeconomic outcomes since the financial crisis," *NBER Macroeconomics Annual*, 31, 283–357.

Christiano, L. J., M. Eichenbaum, and C. L. Evans (1999): "Monetary policy shocks: What have we learned and to what end?" in *Handbook of Macroeconomics*, ed. by J. B. Taylor and M. Woodford, Elsevier, vol. 1 of *Handbook of Macroeconomics*, chap. 2, 65–148.

Clarida, R. H. (2019): "The Federal Reserve's Review of Its Monetary Policy Strategy, Tools, and Communication Practices," Speech, At "A Hot Economy: Sustainability and Trade-Offs," a Fed Listens event sponsored by the Federal Reserve Bank of San Francisco, San Francisco, California.

Clarida, R. H. and D. Coyle (1984): "Conditional Projection by Means of Kalman Filtering," NBER Technical Working Papers 0036, National Bureau of Economic Research, Inc.

Clark, T. E. and M. W. McCracken (2014): "Evaluating Conditional Forecasts from Vector Autoregressions," Working Paper 1413, Federal Reserve Bank of Cleveland.

Conti, A. M., A. Nobili, F. M. Signoretti, et al. (2018): "Bank capital constraints, lending supply and economic activity," Tech. rep., Bank of Italy, Economic Research and International Relations Area.

Del Negro, M., M. P. Giannoni, and C. Patterson (2012): "The forward guidance puzzle," *FRB of New York Staff Report*.

Del Negro, M. and F. Schorfheide (2013): *DSGE Model-Based Forecasting*, Elsevier, vol. 2A of *Handbook of Economic Forecasting*, chap. 2, 57–140.

Doan, T., R. B. Litterman, and C. A. Sims (1986): "Forecasting and conditional projection using realistic prior distribution," Staff Report 93, Federal Reserve Bank of Minneapolis.

Fernández-Villaverde, J., J. F. Rubio-Ramírez, T. J. Sargent, and M. W. Watson (2007): "ABCs (and Ds) of Understanding VARs," *American Economic Review*, 97, 1021–1026.

Giacomini, R. and G. Ragusa (2014): "Theory-coherent forecasting," *Journal of Econometrics*, 182, 145–155.

Giannone, D., M. Lenza, D. Momferatou, and L. Onorante (2014): "Short-term inflation projections: a Bayesian vector autoregressive approach," *International Journal of Forecasting*, 30, 635–644.

Giannone, D., M. Lenza, H. Pill, and L. Reichlin (2012): "The ECB and the interbank market," *The Economic Journal*, 122, F467–F486.

Giannone, D., M. Lenza, and L. Reichlin (2019): "Money, Credit, Monetary Policy, and the Business Cycle in the Euro Area: What Has Changed Since the Crisis?" *International Journal of Central Banking*, 15, 137–173.

Hamilton, J. D. and A. M. Herrera (2004): "Oil Shocks and Aggregate Macroeconomic Behavior: The Role of Monetary Policy: Comment," *Journal of Money, Credit and Banking*, 36, 265–286.

Jarociński, M. (2010): "Conditional forecasts and uncertainty about forecast revisions in vector autoregressions," *Economics Letters*, 108, 257–259.

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

Kilian, L. and S. Manganelli (2008): "The central banker as a risk manager: Estimating the Federal Reserve's preferences under Greenspan," *Journal of Money, Credit and Banking*, 40, 1103–1129.

Kilian, L. and D. P. Murphy (2012): "Why agnostic sign restrictions are not enough: Understanding the dynamics of oil market VAR models," *Journal of the European Economic Association*, 10, 1166–1188.

Leeper, E. M., C. A. Sims, and T. Zha (1996): "What does monetary policy do?" *Brookings Papers on Economic Activity*, 1996, 1–78.

Leeper, E. M. and T. Zha (2003): "Modest policy interventions," *Journal of Monetary Economics*, 50, 1673–1700.

Lucas, R. E. (1976): "Econometric policy evaluation: A critique," in *Carnegie-Rochester Conference Series on Public Policy*, vol. 1, 19–46.

Ludvigson, S. C., S. Ma, and S. Ng (2017): "Shock Restricted Structural Vector-Autoregressions," NBER Working Papers 23225, National Bureau of Economic Research, Inc.

McCulloch, R. E. (1989): "Local Model Influence," *Journal of the American Statistical Association*, 84, 473–478.

McKay, A., E. Nakamura, and J. Steinsson (2016): "The power of forward guidance revisited," *American Economic Review*, 106, 3133–58.

Penrose, R. (1955): "A generalized inverse for matrices," *Mathematical Proceedings of the Cambridge Philosophical Society*, 51, 406–413.

——— (1956): "On best approximate solutions of linear matrix equations," *Mathematical Proceedings of the Cambridge Philosophical Society*, 52, 17–19.

Ramey, V. A. (2016): "Macroeconomic Shocks and Their Propagation," NBER Working Papers 21978, National Bureau of Economic Research, Inc.

Robertson, J. C., E. W. Tallman, and C. H. Whiteman (2005): "Forecasting Using Relative Entropy," *Journal of Money, Credit and Banking*, 37, 383–401.

Rubio-Ramirez, J. F., D. F. Waggoner, and T. Zha (2010): "Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference," *Review of Economic Studies*, 77, 665–696.

Sims, C. A. (1986): "Are forecasting models usable for policy analysis?" *Quarterly Review*, 2–16.

——— (1993): "A Nine-Variable Probabilistic Macroeconomic Forecasting Model," in *Business Cycles, Indicators and Forecasting*, National Bureau of Economic Research, Inc, NBER Chapters, 179–212.

Smets, F. and R. Wouters (2007): "Shocks and frictions in US business cycles: A Bayesian DSGE approach," *American Economic Review*, 97, 586–606.

Tallman, E. W. and S. Zaman (2018): "Combining Survey Long-Run Forecasts and Nowcasts with BVAR Forecasts Using Relative Entropy," *Federal Reserve Bank of Cleveland Working Papers*.

Waggoner, D. F. and T. Zha (1999): "Conditional Forecasts in Dynamic Multivariate Models," *Review of Economics and Statistics*, 81, 639–651.

Online Appendix for

"Structural Scenario Analysis with SVARs"

Juan Antolín-Díaz
London Business School

Ivan Petrella
University of Warwick

Juan F. Rubio-Ramírez*
Emory University
Federal Reserve Bank of Atlanta

May 21, 2020

Contents

A Vector notation for VAR forecasts **2**

B Multiple Solutions for Conditional Forecasts **3**

C Proof of Proposition 1 **4**

D Details of the algorithm **6**

E Alternative implementation using Kalman filtering **9**

F Details on the applications **12**

G Anticipated Policy Shocks in the Smets-Wouters model **15**

H Comparison with Fed Staff Estimates in May 2006 **19**

I Narrative Evidence of Fed forward guidance **22**

J Additional Figures **31**

*Corresponding author: Juan F. Rubio-Ramírez <juan.rubio-ramirez@emory.edu>, Economics Department, Emory University, Rich Memorial Building, Room 306, Atlanta, Georgia 30322-2240.

A Vector notation for VAR forecasts

Assume that we want to forecast the observables for some period ahead using the VAR in Equation (1) of the paper. Then, we can write

$$\mathbf{y}'_{t+h} = \mathbf{b}'_{t+h} + \sum_{j=1}^{h} \boldsymbol{\varepsilon}'_{t+j}\mathbf{M}_{h-j} \text{ for all } 1 < t < T \text{ and all } h > 0,$$

where

$$\begin{aligned}
\mathbf{b}'_{t+h} &= \mathbf{c}\mathbf{K}_{h-1} + \sum_{\ell=1}^{p} \mathbf{y}'_{t+1-\ell}\mathbf{N}_h^{\ell} \\
\mathbf{K}_0 &= \mathbf{I}_n \\
\mathbf{K}_i &= \mathbf{K}_0 + \sum_{j=1}^{i} \mathbf{K}_{i-j}\mathbf{B}_j \text{ if } i > 0 \\
\mathbf{N}_1^{\ell} &= \mathbf{B}_{\ell} \\
\mathbf{N}_i^{\ell} &= \sum_{j=1}^{i-1} \mathbf{N}_{i-j}^{\ell}\mathbf{B}_j + \mathbf{B}_{i+\ell-1} \text{ if } i > 1 \\
\mathbf{M}_0 &= \mathbf{A}_0^{-1} \\
\mathbf{M}_i &= \sum_{j=1}^{i} \mathbf{M}_{i-j}\mathbf{B}_j \text{ if } i > 0 \\
\mathbf{B}_j &= \mathbf{0}_{n\times n} \text{ if } j > p.
\end{aligned}$$

Then

$$\mathbf{y}'_{t+1,t+h} = \mathbf{b}'_{t+1,t+h} + \boldsymbol{\varepsilon}'_{t+1,t+h}\mathbf{M} \text{ for all } 1 < t < T \text{ and all } h > 0,$$

where $\mathbf{y}'_{t+1,t+h} = \left(\mathbf{y}'_{t+1} \dots \mathbf{y}'_{t+h}\right)$, $\mathbf{b}'_{t+1,t+h} = \left(\mathbf{b}'_{t+1} \dots \mathbf{b}'_{t+h}\right)$, $\boldsymbol{\varepsilon}'_{t+1,t+h} = \left(\boldsymbol{\varepsilon}'_{t+1} \dots \boldsymbol{\varepsilon}'_{t+h}\right)$, and

$$\mathbf{M} = \begin{pmatrix} \mathbf{M}_0 & \mathbf{M}_1 & \dots & \mathbf{M}_{h-1} \\ \mathbf{0} & \mathbf{M}_0 & \dots & \mathbf{M}_{h-2} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{M}_0 \end{pmatrix}.$$

It is easy to see that, given that $\mathbf{B} = \mathbf{A}_+\mathbf{A}_0^{-1}$, and $\mathbf{\Sigma} = (\mathbf{A}_0\mathbf{A}_0')^{-1}$, $\mathbf{M}'\mathbf{M}$ depends only on the reduced-form parameters, $\mathbf{B}$ and $\mathbf{\Sigma}$, even though $\mathbf{M}$ depends on the structural parameters, $\mathbf{A}_0$ and $\mathbf{A}_+$. The vector $\mathbf{b}_{t+1,t+h}'$ depends on the history of observables and the reduced-form parameters.

B Multiple Solutions for Conditional Forecasts

It is useful to examine what the general solution in Equations (10) and (12) imply for the three possible cases ($k < nh$, $k = nh$ and $k > nh$). When $k \leq nh$, the system of Equations (8) and (9) is consistent, and Equations (10) and (11) characterize the solution that minimizes the Frobenius norm of $\boldsymbol{\mu}_\varepsilon$ and $\mathbf{\Psi}_\varepsilon$.[1] Therefore, the Penrose solution is the one that envisages the smallest deviation of the mean and covariance matrix of $\tilde{\boldsymbol{\varepsilon}}_{T+1,T+h}$ from the mean and covariance matrix of $\boldsymbol{\varepsilon}_{T+1,T+h}$. In this case, $\mathbf{D}^\star = \mathbf{D}'\left(\mathbf{D}\mathbf{D}'\right)^{-1}$ and Equations (10) and (12) become[2]

$$\boldsymbol{\mu}_\varepsilon = \mathbf{D}'\left(\mathbf{D}\mathbf{D}'\right)^{-1}\left(\mathbf{f}_{T+1,T+h} - \mathbf{C}\mathbf{b}_{T+1,T+h}\right) \tag{B.1}$$

$$\mathbf{\Sigma}_\varepsilon = \mathbf{D}'\left(\mathbf{D}\mathbf{D}'\right)^{-1}\mathbf{\Omega}_f\left(\mathbf{D}\mathbf{D}'\right)^{-1}\mathbf{D} + \left(\mathbf{I}_{nh} - \mathbf{D}'\left(\mathbf{D}\mathbf{D}'\right)^{-1}\mathbf{D}\right). \tag{B.2}$$

Clearly, if $k = nh$ we have that $\mathbf{D}^\star = \mathbf{D}^{-1}$ and Equations (10) and (12) become

$$\begin{aligned}\boldsymbol{\mu}_\varepsilon &= \mathbf{D}^{-1}\left(\mathbf{f}_{T+1,T+h} - \mathbf{C}\mathbf{b}_{T+1,T+h}\right)\\ \mathbf{\Sigma}_\varepsilon &= \mathbf{D}^{-1}\mathbf{\Omega}_f\left(\mathbf{D}^{-1}\right)'.\end{aligned}$$

When $k > nh$, the system is inconsistent, i.e., there is no solution to the system defined by Equations (8) and (9). In other words, not all the restrictions can be satisfied simultaneously. In this case Expressions (10) and (11) are best approximated solutions (see Penrose, 1956), meaning that they minimize

$$\left\|\mathbf{C}\mathbf{b}_{T+1,T+h} + \mathbf{D}\boldsymbol{\mu}_\varepsilon - \mathbf{f}_{T+1,T+h}\right\| \text{ and } \left\|\mathbf{D}\left(\mathbf{I}_{nh} + \mathbf{\Psi}_\varepsilon\right)\mathbf{D}' - \mathbf{\Omega}_f\right\|,$$

respectively, where we are using the Frobenius norm again.[3] In this case, $\mathbf{D}^\star = \left(\mathbf{D}'\mathbf{D}\right)^{-1}\mathbf{D}'$ and

[1] See Golub and Van Loan (1996), p. 55, for a definition of the Frobenius norm.

[2] It is worth noticing that Equations (10) and (12) are equivalent to the ones provided by Andersson et al. (2010), although they only consider the classic conditional forecasting exercise in which conditions are placed on observables. We thank an anonymous referee for suggesting to us this alternative and more straightforward derivation.

[3] For inconsistent systems the Penrose solution gives the same weight to all the constraints. If a researcher wanted to give different weight to the different restrictions, the solution can be easily amended (see Ben-israel and Greville, 2001, p.117).

Equations (10) and (12) become

$$\begin{aligned}\boldsymbol{\mu}_{\varepsilon} &= \left(\mathbf{D}'\mathbf{D}\right)^{-1}\mathbf{D}'\left(\mathbf{f}_{T+1,T+h}-\mathbf{C}\mathbf{b}_{T+1,T+h}\right)\\ \boldsymbol{\Sigma}_{\varepsilon} &= \left(\mathbf{D}'\mathbf{D}\right)^{-1}\mathbf{D}'\boldsymbol{\Omega}_{f}\mathbf{D}\left(\mathbf{D}'\mathbf{D}\right)^{-1}.\end{aligned}$$

C Proof of Proposition 1

In this Appendix we provide a proof for Proposition 1 in the paper. Let the following minimization problem be written

$$\min_{\boldsymbol{\mu},\boldsymbol{\Sigma}} D_{KL}\left(\mathcal{N}\left(\boldsymbol{\mu},\boldsymbol{\Sigma}\right)||\mathcal{N}_{UF}\right)$$

subject to

$$\mathbf{C}\boldsymbol{\mu} = \mathbf{f}_{T+1,T+h} \tag{C.1}$$

$$\mathbf{C}\boldsymbol{\Sigma}\mathbf{C}' = \boldsymbol{\Omega}_{f}. \tag{C.2}$$

For the multivariate Normal case $D_{KL}\left(\mathcal{N}\left(\boldsymbol{\mu},\boldsymbol{\Sigma}\right)||\mathcal{N}_{UF}\right)$ equals

$$\frac{1}{2}\left\{\mathtt{tr}\left[\left(\mathbf{M}'\mathbf{M}\right)^{-1}\boldsymbol{\Sigma}\right]+\left(\mathbf{b}_{T+1,T+h}-\boldsymbol{\mu}\right)'\left(\mathbf{M}'\mathbf{M}\right)^{-1}\left(\mathbf{b}_{T+1,T+h}-\boldsymbol{\mu}\right)-nh+\ln\left[\frac{\det\left(\mathbf{M}'\mathbf{M}\right)}{\det\left(\boldsymbol{\Sigma}\right)}\right]\right\}$$

Solution for the mean The solution for the mean of the conditional forecast solves the following FOCs

$$\mathbf{0}_{k\times 1} = \mathbf{C}\boldsymbol{\mu}-\mathbf{f}_{T+1,T+h}, \tag{C.3}$$

$$\mathbf{0}_{1\times nh} = \left(\mathbf{b}_{T+1,T+h}-\boldsymbol{\mu}\right)'\left(\mathbf{M}'\mathbf{M}\right)^{-1}-\boldsymbol{\lambda}\mathbf{C} \tag{C.4}$$

where $\boldsymbol{\lambda}$ is a $1\times k$ vector of Lagrangian multipliers associated with the mean constraints in (C.1). Equation (C.4) implies that

$$\begin{aligned}\boldsymbol{\lambda}\mathbf{C} &= \left(\mathbf{b}_{T+1,T+h}-\boldsymbol{\mu}\right)'\left(\mathbf{M}'\mathbf{M}\right)^{-1}\\ \boldsymbol{\lambda}\mathbf{C}\mathbf{M}' &= \left(\mathbf{b}_{T+1,T+h}-\boldsymbol{\mu}\right)'\mathbf{M}^{-1}\\ \boldsymbol{\lambda}\mathbf{D} &= \left(\mathbf{b}_{T+1,T+h}-\boldsymbol{\mu}\right)'\mathbf{M}^{-1}\\ \mathbf{0}_{1\times nh} &= \left(\mathbf{b}_{T+1,T+h}-\boldsymbol{\mu}\right)'\mathbf{M}^{-1}\widehat{\mathbf{D}}\end{aligned}$$

where $\widehat{\mathbf{D}}$ denotes the $(nh \times nh)$ annihilator matrix such that $\mathbf{D}\widehat{\mathbf{D}} = \mathbf{0}_{k \times nh}$; therefore, $\widehat{\mathbf{D}} = \mathbf{I}_{nh} - \mathbf{D}'\left(\mathbf{D}\mathbf{D}'\right)^{-1}\mathbf{D} = \mathbf{I}_{nh} - \mathbf{D}^{\star}\mathbf{D}$, where $\mathbf{D}^{\star}$ denotes the Moore-Penrose inverse of $\mathbf{D}$ and has the following properties $\widehat{\mathbf{D}}' = \widehat{\mathbf{D}}$ (since $(\mathbf{D}^{\star}\mathbf{D})' = \mathbf{D}^{\star}\mathbf{D}$) and $\widehat{\mathbf{D}}^{-1} = \widehat{\mathbf{D}}$. Therefore

$$\widehat{\mathbf{D}}\left(\mathbf{M}'\right)^{-1}\boldsymbol{\mu} = \widehat{\mathbf{D}}\left(\mathbf{M}'\right)^{-1}\mathbf{b}_{T+1,T+h}$$

substituting for the definition of the annihilator matrix the previous equation implies that

$$\begin{aligned}
\left(\mathbf{M}'\right)^{-1}\boldsymbol{\mu} &= \left(\mathbf{M}'\right)^{-1}\mathbf{b}_{T+1,T+h} + \mathbf{D}^{\star}\mathbf{D}\left(\mathbf{M}'\right)^{-1}\left(\boldsymbol{\mu} - \mathbf{b}_{T+1,T+h}\right) \\
\boldsymbol{\mu} &= \mathbf{b}_{T+1,T+h} + \mathbf{M}'\mathbf{D}^{\star}\mathbf{D}\left(\mathbf{M}'\right)^{-1}\left(\boldsymbol{\mu} - \mathbf{b}_{T+1,T+h}\right) \\
&= \mathbf{b}_{T+1,T+h} + \mathbf{M}'\mathbf{D}^{\star}\left(\mathbf{C}\boldsymbol{\mu} - \mathbf{C}\mathbf{b}_{T+1,T+h}\right) \\
&= \mathbf{b}_{T+1,T+h} + \mathbf{M}'\mathbf{D}^{\star}\left(\mathbf{f}_{T+1,T+h} - \mathbf{C}\mathbf{b}_{T+1,T+h}\right)
\end{aligned}$$

which is equivalent to the solution for the mean of the conditional forecast in Equation (13) in Section 2 of the paper.

Solution for the variance The solution for the variance of the conditional forecast solves the following FOCs

$$\mathbf{0}_{k \times k} = \mathbf{C}\boldsymbol{\Sigma}\mathbf{C}' - \boldsymbol{\Omega}_f, \tag{C.5}$$

$$\mathbf{0}_{nh \times nh} = \frac{1}{2}\frac{\partial \mathtt{tr}\left[\left(\mathbf{M}'\mathbf{M}\right)^{-1}\boldsymbol{\Sigma}\right]}{\partial \boldsymbol{\Sigma}} - \frac{1}{2}\frac{\partial \ln\left[\det\left(\boldsymbol{\Sigma}\right)\right]}{\partial \boldsymbol{\Sigma}} + \frac{\partial \mathtt{tr}\left[\mathbf{S}\left(\mathbf{C}\boldsymbol{\Sigma}\mathbf{C}' - \boldsymbol{\Omega}_f\right)\right]}{\partial \boldsymbol{\Sigma}}. \tag{C.6}$$

where $\mathbf{S}$ is a square positive semidefinite matrix of Lagrangian multipliers associated with the variance constraints in (C.2). Equation (C.6) implies that

$$\begin{aligned}
\left(\mathbf{M}'\mathbf{M}\right)^{-1} - \boldsymbol{\Sigma}^{-1} + 2\mathbf{C}'\mathbf{S}'\mathbf{C} &= \mathbf{0}_{nh \times nh} \\
\mathbf{I}_{nh} - \mathbf{M}\boldsymbol{\Sigma}^{-1}\mathbf{M}' + 2\mathbf{D}'\mathbf{S}'\mathbf{D} &= \mathbf{0}_{nh \times nh}
\end{aligned}$$

where we have made use of the definition of $\mathbf{D}$ $(= \mathbf{C}\mathbf{M}')$, therefore

$$\mathbf{D}'\mathbf{S}'\mathbf{D} = \frac{1}{2}\left(\mathbf{M}\boldsymbol{\Sigma}^{-1}\mathbf{M}' - \mathbf{I}_{nh}\right) \tag{C.7}$$

Equation (C.7) together with Equation (C.5) implies that

$$\begin{aligned}
\mathbf{D}'\mathbf{S}'\mathbf{D} &= \frac{1}{2}\left(\mathbf{M}\boldsymbol{\Sigma}^{-1}\mathbf{M}' - \mathbf{I}_{nh}\right) \\
\mathbf{S}' &= \frac{1}{2}\mathbf{D}^{\star\prime}\left(\mathbf{M}\boldsymbol{\Sigma}^{-1}\mathbf{M}' - \mathbf{I}_{nh}\right)\mathbf{D}^{\star} \\
&= \frac{1}{2}\mathbf{D}^{\star\prime}\left[\mathbf{M}\left(\mathbf{C}^{*}\boldsymbol{\Omega}_f\mathbf{C}^{*\prime}\right)^{-1}\mathbf{M}' - \mathbf{I}_{nh}\right]\mathbf{D}^{\star} \\
&= \frac{1}{2}\mathbf{D}^{\star\prime}\left[\mathbf{M}\mathbf{C}'\boldsymbol{\Omega}_f^{-1}\mathbf{C}\mathbf{M}' - \mathbf{I}_{nh}\right]\mathbf{D}^{\star} \\
&= \frac{1}{2}\mathbf{D}^{\star\prime}\left[\mathbf{D}'\boldsymbol{\Omega}_f^{-1}\mathbf{D}' - \mathbf{I}_{nh}\right]\mathbf{D}^{\star} \\
&= \frac{1}{2}(\boldsymbol{\Omega}_f^{-1} - \mathbf{D}^{\star\prime}\mathbf{D}^{\star})
\end{aligned}$$

therefore we can solve for the variance under the conditional forecast from Equation (C.7)

$$\begin{aligned}
\mathbf{M}\boldsymbol{\Sigma}^{-1}\mathbf{M}' &= \mathbf{I}_{nh} + \mathbf{D}'(\boldsymbol{\Omega}_f^{-1} - \mathbf{D}^{\star\prime}\mathbf{D}^{\star})\mathbf{D} \\
\mathbf{M}\boldsymbol{\Sigma}^{-1}\mathbf{M}' &= \mathbf{D}'\boldsymbol{\Omega}_f^{-1}\mathbf{D} + (\mathbf{I}_{nh} - \mathbf{D}'\mathbf{D}^{\star\prime}\mathbf{D}^{\star}\mathbf{D}) \\
\boldsymbol{\Sigma}^{-1} &= \mathbf{M}^{-1}(\mathbf{D}'\boldsymbol{\Omega}_f^{-1}\mathbf{D} + \widehat{\mathbf{D}})\left(\mathbf{M}'\right)^{-1}
\end{aligned}$$

which implies that

$$\begin{aligned}
\boldsymbol{\Sigma} &= \mathbf{M}'(\mathbf{D}'\boldsymbol{\Omega}_f^{-1}\mathbf{D} + \widehat{\mathbf{D}})^{-1}\mathbf{M} \\
&= \mathbf{M}'(\mathbf{D}^{\star}\boldsymbol{\Omega}_f\mathbf{D}^{\star\prime} + \widehat{\mathbf{D}})\mathbf{M}
\end{aligned}$$

and is equivalent to the solution for the conditional forecast in Equation (14) in Section 2 of the paper.[4]

D Details of the algorithm

In this section we develop algorithms to implement structural scenario analysis. Conditional-on-variables and conditional-on-shocks forecasts can be implemented as special cases of the former. The algorithms we present can be easily extended to any identification scheme provided that an

[4] To solve for the variance we have made use of the property that given two square matrices $\mathbf{A}$ and $\mathbf{B}$, $(\mathbf{A}+\mathbf{B})^{-1} = (\mathbf{A}^{-1}+\mathbf{B}^{-1})$ if $\mathbf{A}\mathbf{B}^{-1}\mathbf{A} = \mathbf{B}\mathbf{A}^{-1}\mathbf{B}$ (see, e.g., Searle, 1982). The latter conditions in our case are verified: $\mathbf{D}'\Omega_f^{-1}\mathbf{D}\widehat{\mathbf{D}}^{-1}\mathbf{D}'\Omega_f^{-1}\mathbf{D} = \widehat{\mathbf{D}}\mathbf{D}^{\star\prime}\Omega_f\mathbf{D}^{\star}\widehat{\mathbf{D}} = \mathbf{0}_{nh\times nh}$, since from the definition of the annihilator matrix $\mathbf{D}\widehat{\mathbf{D}} = \mathbf{0}$, and $\widehat{\mathbf{D}} = \widehat{\mathbf{D}}'$ and $\widehat{\mathbf{D}}^{-1} = \widehat{\mathbf{D}}$.

algorithm is available to draw from the structural parameters. To simplify both the notation and the exposition, when presenting the algorithm we focus on sign restrictions. Specifically, we extend the Gibbs sampler algorithm in Waggoner and Zha (1999) to implement the structural scenario analysis in set and partially identified SVARs.

In particular, we assume that if $(\mathbf{A}_0, \mathbf{A}_+)$ satisfy the s sign restrictions, then

$$\mathbf{F}(\mathbf{A}_0, \mathbf{A}_+) > \mathbf{0}_{s\times 1}.$$

Before describing the algorithm, it should be clear that our objective is to draw from the following joint posterior

$$p(\tilde{\mathbf{y}}_{T+1,T+h}, \mathbf{A}_0, \mathbf{A}_+ | \mathbf{y}^T, \mathbf{IR}\left(\mathbf{A}_0, \mathbf{A}_+\right), \boldsymbol{R}\left(\tilde{\mathbf{y}}_{T+1,T+h}, \mathbf{A}_0, \mathbf{A}_+\right)), \tag{D.1}$$

where $\mathbf{IR}\left(\mathbf{A}_0, \mathbf{A}_+\right) = \{(\mathbf{A}_0, \mathbf{A}_+) : \mathbf{F}(\mathbf{A}_0, \mathbf{A}_+) > \mathbf{0}_{s\times 1}\}$ are the identification restrictions and $\mathbf{R}(\tilde{\mathbf{y}}_{T+1,T+h}, \mathbf{A}_0, \mathbf{A}_+) = \{\tilde{\mathbf{y}}_{T+1,T+h} : \hat{\mathbf{C}}\tilde{\mathbf{y}}_{T+1,T+h} \sim \mathcal{N}\left(\hat{\mathbf{f}}_{T+1,T+h}, \hat{\mathbf{\Omega}}_f\right)\}$ are the restrictions associated with the structural scenario. Note that $\mathbf{R}(\tilde{\mathbf{y}}_{T+1,T+h}, \mathbf{A}_0, \mathbf{A}_+)$ depends on $(\mathbf{A}_0, \mathbf{A}_+)$ because $\hat{\mathbf{f}}_{T+1,T+h}$, $\hat{\mathbf{C}}$, and $\hat{\mathbf{\Omega}}_f$ depend on $(\mathbf{A}_0, \mathbf{A}_+)$.

As explained by Waggoner and Zha (1999), if one wants to take into account identification and parameter uncertainty, as well as the uncertainty associated with the future path of the shocks and, possibly, the conditional path of the restricted forecast, drawing the posterior distribution described by Equation (D.1) becomes a challenging task. It is tempting, in a first step to, draw the structural parameters from their distribution conditional on $\mathbf{y}^T$ and the identification restrictions, and in a second step to draw $\tilde{\mathbf{y}}_{T+1,T+h}$ conditional on $\mathbf{y}^T$, the restrictions in Equation (15) of the paper, and the structural parameters, using Equations (13) and (14) of the paper. However, this procedure ignores the restrictions in Equation (15) when drawing the structural parameters and, hence, would not lead to a draw from the desired joint posterior described in Equation (D.1). Instead, to draw from the joint distribution of interest, a Gibbs sampler procedure can be constructed that iterates between draws from the described conditional distributions of the structural parameters and $\tilde{\mathbf{y}}_{T+1,T+h}$.[5]

In this appendix we present a particular implementation of Algorithm 1 that follows Faust

[5]The algorithm below is an extension of the one developed by Waggoner and Zha (1999), with the key difference being that we now need to identify the SVAR.

(1998), Uhlig (2005), and Canova and Nicolo (2002), and Rubio-Ramirez et al. (2010) and draw from a conjugate uniform-normal-inverse-Wishart posterior over the orthogonal reduced-form parameterization $(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q})$ and transform the draws into the structural parameterization $(\mathbf{A}_0, \mathbf{A}_+)$. Arias et al. (2018) show that this transformation induces a normal-generalized-normal posterior over the structural parameterization, and highlight the useful properties of using conjugate prior distributions. On the contrary, Baumeister and Hamilton (2015) prefer to use priors stated directly over the structural parameterization, $(\mathbf{A}_0, \mathbf{A}_+)$. Their argument is that since the implementation of Arias et al. (2018) amounts to an implicit and informative prior over the objects of interest to researchers, it is better to define it directly. Baumeister and Hamilton's (2015) approach has an advantage when at least some, if not all, the parameters have an economic interpretation, but could be computationally challenging, especially in large systems. If one wanted to follow that route, our algorithm could be easily modified to draw directly from $(\mathbf{A}_0, \mathbf{A}_+)$.

Algorithm 1. *Initialize* $\mathbf{y}^{T+h,(0)} = [\mathbf{y}^T, \tilde{\mathbf{y}}^{(0)}_{T+1,T+h}]$.

1. *Conditioning on* $\mathbf{y}^{T+h,(i-1)} = [\mathbf{y}^T, \tilde{\mathbf{y}}^{(i-1)}_{T+1,T+h}]$, *draw* $(\mathbf{B}^{(i)}, \mathbf{\Sigma}^{(i)})$ *from the posterior distribution of the reduced-form parameters.*

2. *Draw* $\mathbf{Q}^{(i)}$ *independently from the uniform distribution over the set of orthogonal matrices.*

3. *Keep* $(\mathbf{B}^{(i)}, \mathbf{\Sigma}^{(i)}, \mathbf{Q}^{(i)})$ *if* $\mathbf{F}(f_h^{-1}(\mathbf{B}^{(i)}, \mathbf{\Sigma}^{(i)}, \mathbf{Q}^{(i)})) > \mathbf{0}_{s\times 1}$, *otherwise return to Step 1.*

4. *Conditioning on* $(\mathbf{A}_0^{(i)}, \mathbf{A}_+^{(i)}) = f_h(\mathbf{B}^{(i)}, \mathbf{\Sigma}^{(i)}, \mathbf{Q}^{(i)})$ *and* $\mathbf{y}^T$, *draw* $\tilde{\mathbf{y}}^{(i)}_{T+1,T+h} \sim \mathcal{N}(\boldsymbol{\mu}_y, \mathbf{\Sigma}_y)$ *using Equations (10) and (12) of the paper.*

5. *Return to Step 1 until the required number of draws has been obtained.*

The natural initialization of $\tilde{\mathbf{y}}^{(0)}_{T+1,T+h}$ can be done at the mean of the unconditional forecast or a random draw from its posterior. Steps 1 and 2 draw from posterior uniform-normal-inverse-Wishart over the orthogonal reduced-form parameterization conditional on $\mathbf{y}^{T+h,(i-1)}$. Step 3 checks the sign restrictions, while Step 4 draws $\tilde{\mathbf{y}}^{(i)}_{T+1,T+h} \sim \mathcal{N}(\boldsymbol{\mu}_y, \mathbf{\Sigma}_y)$ using Equations (10) and (12) of the paper. It is important to notice that because Steps 1 and 2 are conditional on $\tilde{\mathbf{y}}^{T+h,(i-1)}$, the restriction in Equation (15) is taken into account when drawing the structural parameters.

Clearly Steps 1-3 are an accept-reject algorithm to draw from the posterior of $(\mathbf{A}_0^{(i)}, \mathbf{A}_+^{(i)})$ conditional on the identification restrictions, as in Rubio-Ramirez et al. (2010), but taking into account the restrictions on $\mathbf{y}^{T+h,(i-1)}$. Step 4 draws $\tilde{\mathbf{y}}^{(i)}_{T+1,T+h}$ conditional on $(\mathbf{A}_0^{(i)}, \mathbf{A}_+^{(i)})$ using

Equation (15). For the case in which the model is exactly identified, such as in the traditional recursive ordering using the Cholesky decomposition (see Kilian and Lutkepohl, 2017), the algorithm is trivially modified by choosing a function h that reflects the restrictions, setting $\mathbf{Q} = \mathbf{I}_n$ in Step 2 and skipping Step 3. The algorithm can also be extended to implement the recently proposed narrative sign restrictions as in Antolin-Diaz and Rubio-Ramirez (2018) and zero restrictions by the methods described in Arias et al. (2018). In both cases, our algorithm requires an importance-sampling step to draw from the posterior of $(\mathbf{A}_0, \mathbf{A}_+)$, which needs to be done after Step 3. This approach is known as Gibbs sampling by sampling-importance-resampling as described in Koch (2007). In particular, the resampling step requires taking multiple draws of the triplet $(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q})$ by repeating Steps 1-3 above, which are then re-weighted using importance weights as described in Antolin-Diaz and Rubio-Ramirez (2018) and Arias et al. (2018) respectively. Finally, a random draw from the reweighted triplets $(\mathbf{B}, \mathbf{\Sigma}, \mathbf{Q})$ is selected before proceeding to Step 4.

E Alternative implementation using Kalman filtering

Clarida and Coyle (1984) show that one can implement the Doan et al. (1986) conditional point forecasts using the Kalman filter. Building on this insight, Banbura et al. (2015) suggest that state-space methods can also be used to obtain conditional-on-observables forecasts as in Section 2. We now show that this equivalent algorithm using the Kalman filter can also be applicable to the structural scenario analysis as well. As we will see, however, the Kalman filter implementation will be in most cases computationally less efficient than the general, closed-form solutions provided in the paper. Moreover, it does not easily lend itself to implement more flexible linear restrictions such as the "Back to the Price Level Target" exercise described in Section 5.

The state-space representation

We begin by outlining the state-space representation that allows the use of the Kalman filter to compute conditional forecasts. Consider the following state-space form consistent with the model in Section 2:

$$\mathbf{z}_t' = \mathbf{D} + \mathbf{z}_{t-1}'\mathbf{F} + \boldsymbol{\varepsilon}_t'\mathbf{G} \tag{E.1}$$

$$\tilde{\mathbf{y}}_t' = \mathbf{H}\mathbf{z}_t' \tag{E.2}$$

The state vector $\mathbf{z}_t' = [\mathbf{y}_t', \ldots, \mathbf{y}_{t-p+1}', \boldsymbol{\varepsilon}_t']$ stacks the endogenous variables, their lags, and the exogenous shocks, where $\mathbf{y}_t$ is an $n \times 1$ vector of variables at time t, and $\boldsymbol{\varepsilon}_t$ is an $n \times 1$ vector of structural shocks at time t. The measurement vector $\tilde{\mathbf{y}}_t' = [\mathbf{y}_t', \boldsymbol{\varepsilon}_t']$ stacks the endogenous variables and the exogenous shocks. As we will see below, since the structural shocks are unobserved, the measurement vector will contain missing observations corresponding to the entries of $\boldsymbol{\varepsilon}_t$. The system matrices $\mathbf{D}$, $\mathbf{F}$, $\mathbf{G}$, and $\mathbf{H}$ are constructed from the SVAR parameters as shown below:

$$\mathbf{D} = \begin{bmatrix} \mathbf{c} \\ \underset{n\cdot p\times 1}{\mathbf{0}} \end{bmatrix}, \qquad \mathbf{F} = \begin{bmatrix} \underset{n\cdot p\times n\cdot p}{\tilde{\mathbf{B}}} & \underset{n\cdot p\times n}{\mathbf{0}} \\ \underset{n\times n\cdot p}{\mathbf{0}} & \underset{n\times n}{\mathbf{0}} \end{bmatrix}, \qquad \mathbf{G} = \begin{bmatrix} h(\boldsymbol{\Sigma})\mathbf{Q}^{-1} \\ \underset{n\cdot(p-1)\times n}{\mathbf{0}} \\ \underset{n\times n}{\mathbf{I}} \end{bmatrix}, \qquad \mathbf{H} = \begin{bmatrix} \underset{n\times n}{\mathbf{I}} & \underset{n\times n\cdot p}{\mathbf{0}} \\ \underset{n\times n\cdot p}{\mathbf{0}} & \underset{n\times n}{\mathbf{I}} \end{bmatrix},$$

where $\tilde{\mathbf{B}}$ is known as the companion matrix of the VAR and is defined as

$$\tilde{\mathbf{B}} = \begin{bmatrix} \mathbf{B}_1 & \mathbf{B}_2 & & \ldots & \mathbf{B}_p \\ \mathbf{I} & \mathbf{0} & & & \\ \mathbf{0} & \mathbf{I} & \ddots & & \vdots \\ \vdots & & & \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \ldots & \ldots & \mathbf{0} & \mathbf{I} \end{bmatrix}.$$

The SVAR representation in Equation 1 of the paper and the state-space representation in Equations (E.1)-(E.2) described above are equivalent, and both can be used to compute conditional forecasts. With the model in state-space one can use the simulation smoother (see, e.g., Durbin and Koopman, 2012) to generate a draw of the state vector $\tilde{\mathbf{y}}_t'$, conditional on the observations and the parameters of the model. Therefore, as outlined by Banbura et al. (2015) for the reduced-form case, one can treat the future path of conditioning variables as observed, while leaving the future path of the remaining variables as missing observations.[6]

The same procedure can be adapted for the case of the structural scenario. Suppose that one

[6] To allow for missing observations, the standard Kalman filtering recursions must be modified. For this purpose, a number of equivalent solutions have been proposed (see Durbin and Koopman (2012), p. 112, for a textbook treatment). Here we follow the approach of Mariano and Murasawa (2003) and Camacho and Perez-Quiros (2010), which essentially impose the condition that the Kalman gain associated with the missing observation is mechanically zero. This implies that the Kalman filter "skips" missing observations, effectively marginalizing the likelihood with respect to the missing observations (see Brockwell and Davis (1991), Section 12.3, and Brockwell et al. (1990)).

wants to produce forecasts for h time periods into the future. Denote with $\star$ an observation for which all data entries are missing, denote $(\hat{y}_{i,t}, \star)'$ for $T+1 \leq t \leq T+h$ a post-sample observation in which the value of endogenous variable i is assumed to be known and fixed, the rest being missing data, and denote $(\hat{\varepsilon}_{j,t}, \star)'$ for $T+1 \leq t \leq T+h$ a post-sample observation in which the value of structural shocks j is constrained (for instance to reflect its unconditional standard normal distribution), the observation for other shocks being missing data. All of the cases mentioned above can be considered special cases of the state-space representation (E.1)-(E.2), in which the measurement vector is constructed as $\tilde{\mathbf{y}}_t'$ in the following way:

$$\tilde{\mathbf{y}}_t' = [\mathbf{y}_t', *], \qquad \text{for } 1 \leq t \leq T, \tag{E.3}$$

$$\tilde{\mathbf{y}}_t' = [(\hat{y}_{i,t}, *)', (*, \theta_t)'], \qquad \text{for } T+1 \leq t \leq T+h, \tag{E.4}$$

where θ_t denotes the vector of constrained structural shocks and is a $1 \times n-1$ draw from the unconditional distribution of ε_t, i.e., from an independent standard normal distribution. In other words, n columns and h rows of missing observations are appended to the data set containing the endogenous variables, the known future path for variable i for periods $T+1$ to $T+h$ is used to fill the corresponding missing observations, and the missing observations for the structural shocks *other than* the shock of interest are filled in with draws from their unconditional distribution.[7] In Step 4 of Algorithm 1, one can then draw $y_{T+1,T+h}^{(i)}$ using the Kalman smoother, as opposed to a multivariate Gaussian distribution with mean and variance given by Equations (10)-(12) of the paper. These two procedures produce identical results.

Computational efficiency

Since the Kalman filtering-based procedure provides results identical to those of the algorithm outlined in the paper, we now assess which of the two methods is more computationally efficient. Equations (10)-(12) of the paper involve operations of high dimensional matrix objects which modern computing langues are well suited to compute efficiently, whereas the procedure based on the Kalman filter works recursively leading to a higher computational burden.

Table E.1 investigates which of the two procedures is more efficient numerically for different

[7] If one is interested in computing only point forecasts, filling them with zeros, the unconditional mean, is appropriate.

Table E.1: Computational efficiency of conditional forecast methods

		$p=1$		$p=6$		$p=12$		$p=24$	
	Horizon	Baseline	KF	Baseline	KF	Baseline	KF	Baseline	KF
n=3	h=1	**0.49**	27.37	**0.5**	34.68	**0.52**	51.77	**0.56**	88.15
	h=6	**0.8**	28.09	**1.12**	35.32	**2.25**	49.02	**2.11**	90.35
	h=12	**1.29**	28.48	**3.1**	36.17	**5.53**	50.24	**5.42**	95.68
	h=24	**2.78**	29.76	**5.06**	37.92	**8.89**	52.72	**20.28**	98.19
	h=60	**12.37**	33.61	**18.51**	43.38	**31.35**	59.67	**69.64**	113.28
n=7	h=1	**0.53**	35.57	**0.55**	66.81	**0.51**	141.74	**0.61**	368.3
	h=6	**1.42**	36.21	**2.47**	67.34	**3.07**	139.76	**3.03**	375.16
	h=12	**2.81**	37.17	**4**	69.25	**5.43**	147.83	**7.79**	386.21
	h=24	**9.16**	38.68	**11.81**	72.16	**16.72**	148.68	**28.05**	403.88
	h=60	77.07	**43.67**	**83.9**	84.62	**102.7**	174.42	**150.89**	458.84
n=11	h=1	**0.58**	34.04	**0.61**	93.49	**0.55**	200.4	**0.68**	697.13
	h=6	**2.12**	35.06	**2.58**	91.7	**3**	209.97	**3.82**	738.55
	h=12	**5.61**	36.02	**7.6**	94.64	**9.97**	218.71	**11.62**	751.22
	h=24	**25.31**	38.38	**25.36**	99.79	**31.52**	231.55	**48.95**	821.48
	h=60	210.07	**44.63**	216.67	**126.62**	**247.11**	277.71	**308.01**	964.74
n=27	h=1	**0.9**	131.52	**0.65**	579.93	**1.28**	2241.43	**0.99**	12193.3
	h=6	**8.98**	130.98	**9.17**	622.57	**10.34**	2314.67	**12.88**	12335.05
	h=12	**35.94**	135.65	**40.92**	599.76	**45.23**	2351.54	**52.18**	12644.66
	h=24	224.8	**135.44**	**224.38**	625.57	**248.64**	2461.27	**285.89**	13288.57
	h=60	6301.38	**158.01**	3874.75	**732.05**	3980.7	**2857.22**	**4089.79**	15290.47

Note: The numbers denote the average time, in seconds (over 100 iterations), taken to calculate 1000 draws of conditional forecasts. n denotes the number of variables in the VAR, h denotes the forecast horizon and p denotes the number of lags in the VAR. The time of the fastest algorithm in each case is marked in boldface. These computations were performed in Matlab 2015a on an HP Z230 SFF workstation with an Intel Core i7-4770 CPU 3.40GHz processor and 7.74GB of usable RAM.

hypothetical problem settings.[8] The baseline implementation of Algorithm 1 is faster in the majority of cases, with the gains particularly large whenever the number of variables n or lags p is large. Only whenever h increases a lot relative to the number of variables and lags is small, drawing the conditional forecast using the Kalman smoother is more efficient.

F Details on the applications

Application 1: Monetary policy alternatives

Data

We work with the data set of Smets and Wouters (2007), which contains seven key US macroeconomic time series at the quarterly frequency: real output, consumption, investment, wages and hours

[8] We use deJong's (1988) implementation of the Kalman filter, as Banbura et al. (2015) show that this implementation is computationally faster than alternative implementations.

worked per capita, inflation as measured by the GDP deflator and the federal funds rate.[9] Their data set starts in Q1 1966, and we update it through Q4 2019. For the comparison with the DSGE model in Section 5, we use the original dataset in Smets and Wouters, downloaded from the AER website. We verified that the original dataset corresponds to the vintage available in March 2005. We make one modification to the dataset, which is to substitute the GDP deflator for the PCE deflator excluding food and energy as an indicator for inflation. We use the data for the March 2005 vintage. The motivation for this change is to be consistent with the fact that the Federal Reserve's inflation objective is stated in terms of the PCE deflator, and in particular many of its communications and internal documents refer to the less volatile ex-food and energy measure as an more reliable indicator of future inflation tendencies. Over the sample considered, the core PCE and the overal GDP deflator are 93% correlated.

For the applications updated through 2019, we udpate the dataset using data from Haver Analytics and using the corresponding tickers:

1. Gross Domestic Product (GDV@USECON)
2. Gross Domestic Product - Implicit Price Deflator (JGDP@USECON)
3. Core PCE Deflator (LXEHL@USECON)
4. Personal Consumption Expenditures, Non-Durable Goods and Services (CS@USECON + CN@USECON)
5. Fixed Private Investment + Consumer Durables (F@USECON + CD@USECON)
6. Total Economy Hours Worked (LXEHL@USECON)
7. Compensation of Employees (YCOMP@USECON)
8. Civilian Non-Institutional Population (LNN@USECON)
9. Federal Funds Rate (FFED@USECON)

Output, Consumption and Investment are divided by the GDP deflator and by population to obtain real, per-capita series. Hours are divided by population. Compensation is divided by hours worked and the GDP deflator to obtain real compensation per hour. To construct the measure of expect

[9] See the data appendix in Smets and Wouters (2007) for exact definitions of the variables.

interest rates we use series from the Philadelphia Fed Survey of Professional Forecasters 3 month T-bill expecations:

1. Current Quarter (ASAFC3T@SURVEYS)
2. 1 Quarter Ahead (ASAF13T@SURVEYS)
3. 2 Quarter Ahead (ASAF23T@SURVEYS)
4. 3 Quarter Ahead (ASAF33T@SURVEYS)
5. 4 Quarter Ahead (ASAF43T@SURVEYS)

Then, the variable

$$S_{t,h} \equiv E_t^{SPF}[i_{t+h}] - E_{t-1}^{SPF}[i_{t+h}] \tag{F.1}$$

for $h = 0, \ldots, 3$ measures the *surprise* to expectations, i.e., the change in the information set, from $t-1$ to t. Campbell et al. (2012) point out that surprises across horizons h exhibit a strong factor structure, so we take the average across horizons, $\bar{S}_t \equiv 1/4 \sum_{h=0}^{3} S_{t,h}$ as a measure of the surprises to the average expected interest rate for the next year.

Estimation settings

We estimate a VAR with $p = 4$ lags using Bayesian methods and an informative "Minnesota" prior (Doan et al., 1986; Sims, 1993) over the reduced-form parameters. Also see Karlsson (2013) In order to deal with the absence of data for the survey in the early part of the sample, we need to add an additional step to Algorithm 1. We cast the model in state-space and, conditional on the previous draw of the reduced-form parameters, we draw the missing observations using a simulation smoother (e.g., Durbin and Koopman, 2012)

Application 2: Stress-testing

Data

In this section we report the definition and Haver mnemonics for the data used in Application 2, which are appended to the dataset used in Application 1:

- 3-month Treasury bill rate (FTBS3@USECON)

- unemployment rate (LR@USECON)
- 10-year yield (FCM10@USECON)
- S&P 500 index (SP500@USECON)
- house price index (USRSNHPM@USECON)
- oil price deflated by the core PCE deflator (PZTEXP@USECON/JCXFE@USECON)
- credit spread is the BAA credit spread minus the treasury long-term composite yield (FBAA@USECON-FLTG@USECON)
- TED spread (C111FRED@OECDMEI - FTBS3@USECON)
- ROE for all FDIC Insured Institutions (USARQ@FDIC).

The data are quarterly from 1966 to 2019.

Estimation settings

The VAR is estimated with gour lags and the Minnesota prior implemented using dummy observations (see Del Negro (2011) for details) over the reduced-form parameters.

G Anticipated Policy Shocks in the Smets-Wouters model

In the empirical section we compare the results from the VAR to the ones from a modified version of the Smets and Wouters (SW) model. In this appendix we describe the main modifications of the model and explain the use we make of it, in particular to inform us on the restrictions to impose for the identification of the anticipated monetary policy shock and on the feasibility of the identification of the structural shocks from the VAR.

The only modification of the model is the introduction of anticipated monetary policy shocks. In particular, modify the policy rule in the original model and introduce news shocks as follows

$$r_t = \rho r_{t-1} + (1-\rho)[r_\pi \pi_t + r_Y(y_t - y_t^p)] + \rho_{\Delta y}[(y_t - y_t^p) - (y_{t-1} - y_{t-1}^p)] + \epsilon_t^r + \sum_{j=0}^{3} v_{t-j,j}^r \quad \text{(G.1)}$$

where $v_{t-j,j}^r$ is an anticipated monetary policy shock, which corresponds to deviations from the normal conduct of policy revealed at time $t-j$ yet affecting the interest rate with a delay of j

Figure G.1: Federal Fund Surprises

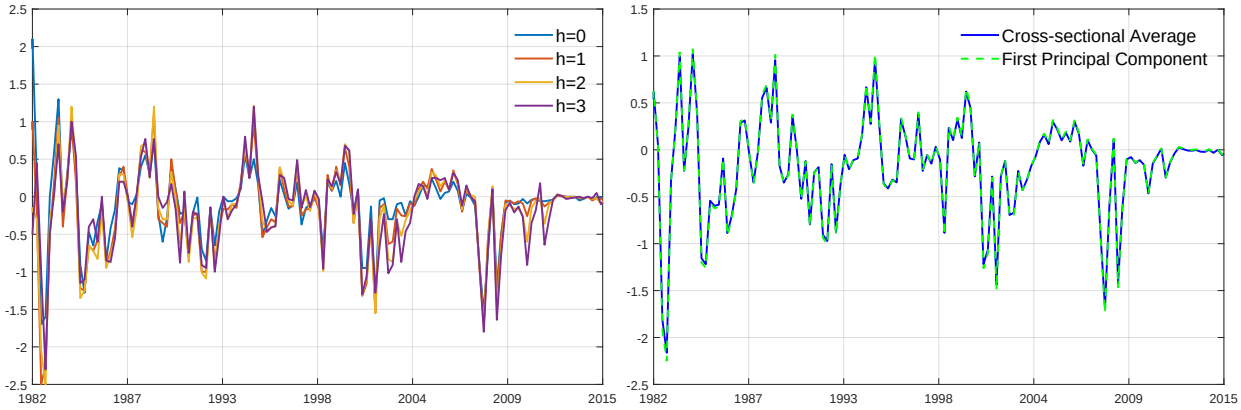

Note. The left chart reports $S_{t,h} \equiv E_t^{SPF}[i_{t+h}] - E_{t-1}^{SPF}[i_{t+h}]$, for $h = 0, \ldots, 3$. The right panel of the figure contrasts the cross sectional average of the surprises with their first principal component.

quarters. In order to capture that idea that forward guidance announcements affect the path of the interest rate over multiple quarters, we assume that news shocks follow a simple MA(4) structure:

$$v_{t-j,j}^r = \eta_{t-j}^r \quad j = 0, \ldots, 3 \tag{G.2}$$

The way we enter anticipated monetary policy shocks into the model is in line with Campbell et al. (2012). The simplifying assumption of a single factor structure in equation G.2 reflects the observations that revisions in interest rate forecast at different quarters display a strong comovement and roughly equal volatility, as highlighted in Figure G.1. Therefore the key difference between a standard (unanticipated) monetary policy shock and a 'forward' guidance shock is that the latter is expected to affect the path of the interest rate for prolonged periods of time. In order to stay as close possible to the original SW model we assume that the law of motion of η_t^r follows an AR(1) and fix the parameters of this to be the same as the one of the unanticipated monetary policy shock.

The first question that we ask from the model is whether a seven variables VAR – featuring the original variables in the SW model – would able to recover an anticipated monetary policy shock, i.e. more specifically whether the VAR is informationally sufficient or fundamental. Fernández-Villaverde et al. (2007) put forward a simple metric, the 'poor man invertibility condition' (PMIC), which we can check on the implied VAR solution in the theoretical model. Fixing the parameters estimates to the modal estimate in SW, the PMIC shows that a 7 variables VAR is not fundamental in the structural shocks. However, expanding the model so as to include the average forecast revisions in

the short term interest rate suggests that a VAR augmented with a surprise variable calculated as described in the text satisfies the invertibility requirement and is therefore informationally sufficient. This is computed in line with the definition given in Section 5 of the paper (equation (F.1) above) and with expectations computed under the assumption of rationality. However, we do not need the survey expectations to be fully rational in order for the VAR to be fundamental and identification to work.

We also use the model to derive plausible restrictions for the identification of the SVAR. Simulations from the model suggest that an anticipated monetary policy shock is associated with a pattern of signs of the IRFs which is consistent with the one of a traditional monetary policy shock. An anticipated monetary policy shock is also associated with a movement in the interest rate surprise variable, $\bar{S}_t$ that goes in the same direction than the shock. When comparing the impact of an unanticipated and an anticipated MP shock it is clear that, for a shock normalized to have the same initial impact on the short term interest rate, the anticipated monetary policy shocks are associated with larger average revisions in $\bar{S}_t$ than in the unanticipated case.[10] Figure G.2 compares the impulse response function of the SVAR and the SW model for a monetary policy (MP) shock and the anticipated monetary policy (AMP) shock. The SVAR model is associated with a more persistent effect of both shocks, but otherwise the IRFs are remarkably similar despite the fact that the shocks in the SVAR are estimated imposing only a limited number of identification restrictions.

[10] Intuitively, with policy rule (G.1) a shock to η_t^r that rises the interest rate by 25bp increases the surprise variable $\bar{S}_t$ roughly 4 times more than a shock to ϵ_t^r.

Figure G.2: Impulse Responses: VAR vs. SW model

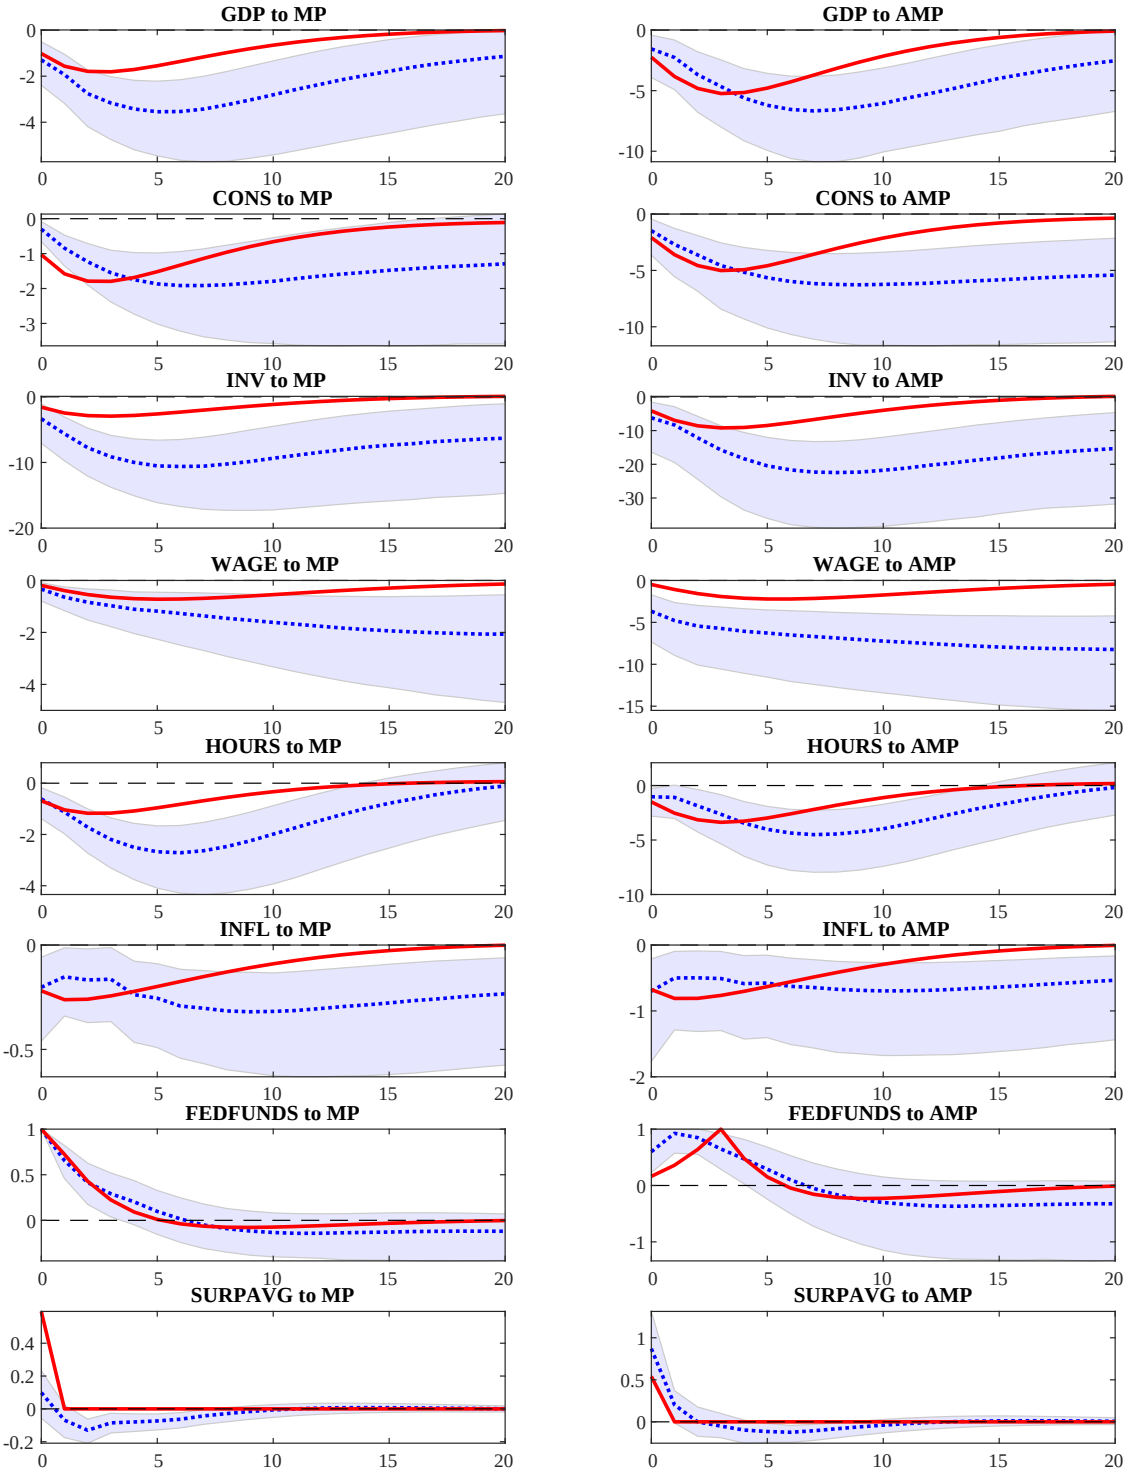

Note. Impulse responses to a monetary policy shock (left column) and an anticipated monetary policy shock (right column). Red continuous line corresponds to the IRFs from the SW model, whereas blue dotted lines corresponds to the IRFs from the SVAR, the blue shaded area represents the 68% (point-wise) HPD credible sets for the IRFs.

H Comparison with Fed Staff Estimates in May 2006

In this section we compare results using our SVAR with internal exercises conducted by the Fed Staff using the FRB/US model and available as part of declassified historical materials. We have taken the results of the exercises published in the Bluebook dated May 5, 2006.[11] We re-estimate the model using the vintage of data available at that time. We also appended the unemployment rate to the dataset used in the VAR in order to obtain a conditional forecast of this variable as well that we can compare with the Fed staff estimates.

In this meeting, the FOMC was presented with two alternatives. The first is the Greenbook baseline in which the fed funds rate is raised to 5% and then lowered below 4.75%. To shed light on the implications of ending the process of policy firming, a second simulation considers a policy trajectory in which the federal funds rate is held constant at its current level through 2010 rather than being raised to 5%. Figure H.1 shows the comparison of the two scenarios, with the corresponding results published by the Fed staff. In the description of the Bluebook, the staff mentions that " monetary policy adjusts to hold core PCE inflation at about 2 percent" so we condition on inflation being at 2% in expectation at the end of the horizon. They also state the "assumption that this policy path is fully anticipated by financial market participants" which in our framework is operationalized as allowing both anticipated and unanticipated shocks to deviate from their unconditional distribution, and "the absence of any significant supply shocks or changes in expectations", which in our setting we model as all non-policy shocks to be at their unconditional distribution.

It is interesting to note that while the median of the SVAR result is quite close to the FRB/US result for the unemployment rate, inflation is projected to be higher with the SVAR than with the FRB/US in both scenarios. This is a result of the federal funds rate staying persistently below the SVAR unconditional forecast throughout the forecast horizon, reflecting a sequence of monetary policy shocks which also drive up inflation above the 2% target. One interpretation of our results is that the SVAR would have suggested that the Fed staff forecasts are too optimistic on the inflation front, and that further interest rate increases would be required to stabilize inflation. In the event, this is what happened, with the FOMC ultimately increasing the federal funds rate to 5.25% and inflation stabilizing around 2%. The mode of the KL is 0.80 for both exercises. The reason the

[11] See https://www.federalreserve.gov/monetarypolicy/files/FOMC20060510bluebook20060505.pdf (Chart 5).

scenario is deemed to be a relatively implausible one is because it embodies a highly persistent deviation from the unconditional scenario at the end of the forecast horizon.

Figure H.1: FOMC Monetary Policy Alternatives: May 2006

(a) Baseline

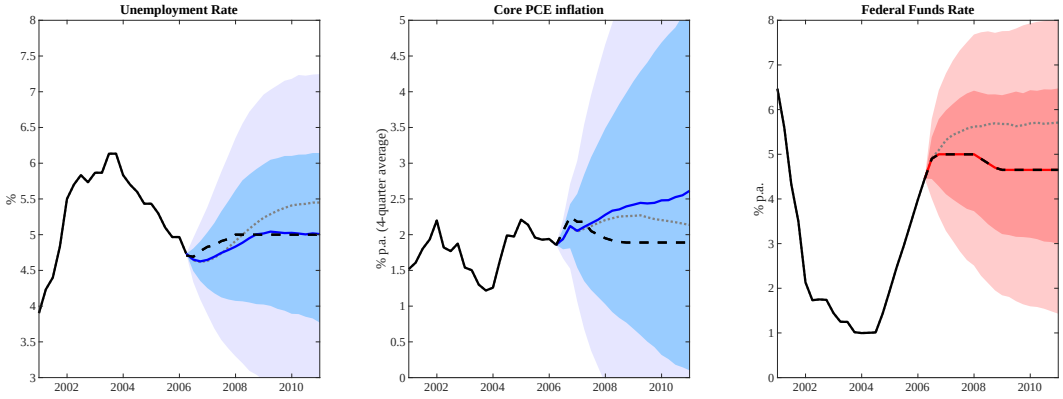

(b) Lower near-term funds rate

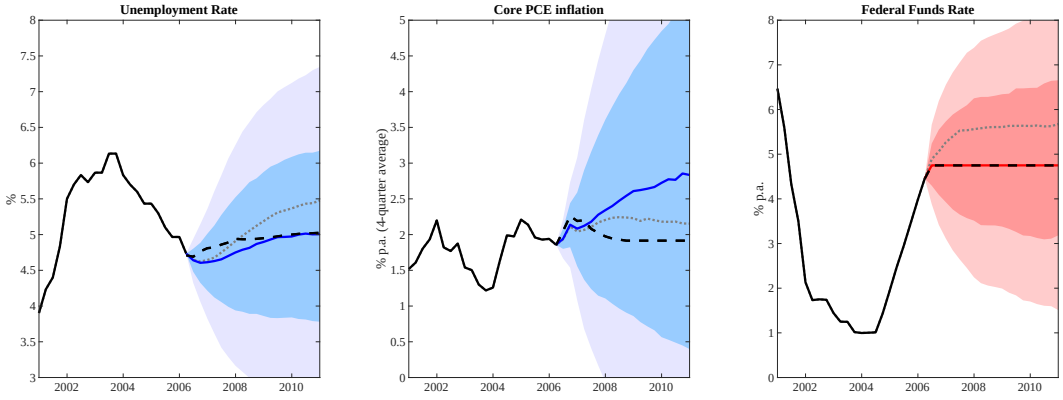

Note: For each panel, the solid black lines represent actual data. In the third column, the solid red line and associated shaded areas are the median, 40 and 68% pointwise credible sets around the conditioning assumptions for the federal funds rate. In the first and second panel, the solid blue line and associated shaded areas are the median, 40 and 68% pointwise credible sets around the forecasts of the other variables given by the SVAR. The dotted gray lines represent the median unconditional forecast from the SVAR. The broken black lines represent instead the results of the simulations from the FRB/US model as reported in the Fed Bluebook in May 2006.

I Narrative Evidence of Fed forward guidance

In this appendix we provide evidence in support of the narrative identification restrictions used to complement the sign restrictions in the paper. In particular we report extract of FOMC documents in the periods 2003-2004 and 2011-2012 that justify our choices.

FOMC communication and forward guidance: 2003-2004

In order to enhance the communications with the public, since January 2000 the FOMC started to issue a brief public statement after every FOMC meeting. The FOMC's immediate announcement has contained a summary rationale for its policy decision and encapsulate the Committee's view of future developments. On some occasions, the announcement has directly hinted at the likelihood of changes in policy. The statement has been kept brief, partly to improve the predictability of the market reaction. Unsurprisingly, from the very beginning the FOMC statements have attracted considerable attention from the media and the markets, with the language closely watched with regard both to the current decision and to clues about future policy settings. In particular, special attention is paid to any changes in wording since the last statement. The first statements used the "balance of risk" assessment as a way of conveying the most likely 'tilt' in future policy. This practice had however at time come under intense scrutiny. In fact, "the meaning of the balance-of risks sentence is somewhat vague, giving rise to differences of opinion even among FOMC members and possibly leading to difficulties among outsiders in understanding its subtleties." (p. 191).[12]

Governor Bernanke highlights the importance of giving more precise guidance on the policy stance. In particular, for the August 12, 2003 meeting transcript Bernanke says (page 131):

"Our discussion yesterday and our experience of the past year have shown that effective communication is a big part of successful monetary policymaking. [...] We can communicate to financial markets how we expect to manage short-term interest rates in the future. If we then follow through on our commitments, the power of our interest rate instrument is multiplied many times over. [...] Nevertheless, communication that is entirely qualitative in nature also risks misunderstanding. [...] In an environment of low inflation and low interest rates, we need to seek ever greater clarity of communication to the markets and to the public." In mid-2003 the FOMC starts to give a more explicit reference to the future monetary policy path.

[12] See https://www.federalreserve.gov/monetarypolicy/files/FOMC20030624memo01.pdf

The September 14, 2003 FOMC meeting is dedicated to the discussion of the FOMC communication policy and procedure around it. In this meeting Gov. Bernanke makes the point that explicit communication in the FOMC statement should be used to convey information about the policy that is not only about the FOMC view on the economy: *"[. . .] we shouldn't be giving unconditional information about our interest rate setting. We should be trying to guide the market by providing conditional forecasts about which direction our policy may be going in light of our objectives, our views, or the nature of our rule."*

August 12, 2003 FOMC Meeting: Negative anticipated monetary policy shock

FOMC Statement *The Committee perceives that the upside and downside risks to the attainment of sustainable growth for the next few quarters are roughly equal. In contrast, the probability, though minor, of an unwelcome fall in inflation exceeds that of a rise in inflation from its already low level. The Committee judges that, on balance, the risk of inflation becoming undesirably low is likely to be the predominant concern for the foreseeable future. In these circumstances, the Committee believes that policy accommodation can be maintained for a considerable period.*

Minutes *While the Committee could not commit itself to a particular policy course over time, many of the members referred to the likelihood that the Committee would want to keep policy accommodative for a longer period than had been the practice in past periods of accelerating economic activity. Reasons for such an approach to policy stemmed from the need to encourage progress toward closing the economy's currently wide output gap and, with inflation already near the low end of what some members regarded as an acceptable range, to resist significant further disinflation. [. . .] At the same time, maintaining an accommodative policy stance was seen as involving little risk of inducing rising inflation so long as high levels of excess capacity and very competitive markets continued to characterize economic conditions.*

Additional support for the choice of this data as an example of forward guidance comes from the Transcripts of the 28 Jan 2004 FOMC meeting. For instance, in his prepared remarks the FOMC secretary Reinhart states: *"At times over the past six months, members have chaffed at the constraint imposed by the commitment to keep policy accommodative for a considerable period. But by being explicit to the public about this self-imposed constraint, you did help limit the tendency of market participants to build in unhelpfully aggressive expectations of policy firming, thereby keeping financial conditions accommodative at a time when you might have been concerned about the efficacy*

of alternative monetary policy actions. In the event, the expansion of aggregate demand did not falter, and there was no need to dig deeper into the toolkit of policymaking."

In the same meeting Gov. Kohn makes the following remark: *"The "considerable period" phrase was inserted as a form of unconventional policy when we were concerned about deflation and the lower nominal bound."*

Jan. 28, 2004 FOMC Meeting: Negative anticipated monetary policy shock

The Committee perceives that the upside and downside risks to the attainment of sustainable growth for the next few quarters are roughly equal. The probability of an unwelcome fall in inflation has diminished in recent months and now appears almost equal to that of a rise in inflation. With inflation quite low and resource use slack, the Committee believes that it can be patient in removing its policy accommodation.

Minutes *"With regard to the wording of the Committee's press statement to be released shortly after the meeting, members discussed at some length the desirability of retaining a reference from earlier statements to the prospect that an accommodative policy could be maintained "for a considerable period." [...] All the members agreed that a change in wording was desirable, not to signal a policy tightening move in the near term, but rather to increase the Committee's flexibility to take such an action when it was deemed to be desirable and to underline that any such decision would be made on the basis of evolving economic conditions. [...] A number of members commented that expectations of sustained policy accommodation appeared to have contributed to valuations in financial markets that left little room for downside risks, and the change in wording might prompt those markets to adjust more appropriately to changing economic circumstances in the future. "*

A reading of the transcripts of the discussion in the FOMC highlights that the change in wording in the statement it transpires that the statement is meant to highlight the FOMC different assessment of state of the economy and therefore to align market expectations to the one of the FOMC, which sees a future tightening of policy more likely than what the markets were predicting.

Transcripts *" That there remains considerable mass on the possibility of earlier tightening embedded in financial market prices implies that there is a potential for the rally to be extended as expectations correct more to the Greenbook baseline. Why market participants might still expect earlier and more significant firming is evident in the bottom right panel: A survey of economists at eight dealers indicates that their outlook for inflation is decidedly less subdued than the staff's, even*

with a growth forecast that is less robust than the staff's. A distinct possibility is that many in the market have a gloomier view of the prospects for the growth of aggregate supply [...]"

"In the Bluebook, we suggested adopting the notion of "patience" that the Chairman introduced in a recent speech, which market participants would probably take as implying that the Committee viewed events as such that it could be gradual in firming policy."

Chairman Greenspan speaking: *"[...] I think today is the day we should adjust our press statement and move to a reference to "patience." I think the downside risks to that change are small. I do think the market will react "negatively" as we used to say, but I'm not sure such a reaction would have negative implications, quite frankly. [...] If we go to "patience," we will have full flexibility to sit for a year or to move in a couple of months. [...] "*

May 4, 2004 FOMC Meeting: Positive anticipated monetary policy shock

"The Committee perceives the upside and downside risks to the attainment of sustainable growth for the next few quarters are roughly equal. Similarly, the risks to the goal of price stability have moved into balance. At this juncture, with inflation low and resource use slack, the Committee believes that policy accommodation can be removed at a pace that is likely to be measured."

The keyword in the statement, as evident from the minutes and traspits of the discussion, is "measured". This is meant to convey the intention of moving toward a tighter policy, but with at a pace that is slower than what market participants are expecting.

Minutes: *The Committee also discussed at length the advantages and disadvantages of modifying or dropping its statement in the announcement following the March meeting that "With inflation quite low and resource use slack, the Committee believes that it can be patient in removing its policy accommodation." All of the members agreed that, with policy tightening likely to begin sooner than previously expected, the reference to patience was no longer warranted. The Committee focused instead on a formulation that would emphasize that policy tightening, once it began, probably could proceed at a pace that would be "measured." A number of policymakers were concerned that such an assertion could unduly constrain future adjustments to the stance of policy should the evidence emerging in coming months suggest that an appreciable firming would be appropriate. Others, however, saw substantial benefits to inclusion of the proposed language. These members noted that current economic circumstances made it likely that the process of returning policy to a more neutral setting would be more gradual, once under way, than in past episodes when inflation was*

well above levels consistent with price stability. In addition, some policymakers observed that the timing and magnitude of future policy adjustments would ultimately be determined by the Committee's interpretation of the incoming data on the economy and prices rather than by its current expectation of those developments. On balance, all the members agreed that they could accept an indication in the statement that ". . . policy accommodation can be removed at a pace that is likely to be measured."

Transcripts *In the Bluebook, we offered modified language that the "Committee believed that policy accommodation can likely be removed at a measured pace" to emphasize that you see yourselves entering a tightening phase but one that likely will not be as aggressive as in prior episodes. That emphasis seemed important because, as in the table at the bottom left, market participants currently anticipate about 200 basis points of firming in the next year. That is probably influenced by the experience of the three prior tightening phases, in which the funds rate moved an average of 2 2/3 percentage points higher within the first year. But with inflation lower now than in those prior episodes, you may see less need at this juncture for such firming, perhaps viewing as more likely something on the order of the 100 basis point increase assumed by the staff. In the Greenbook, the gradual realization by investors that you will not need to tighten as sharply as they now believe imparts an accommodative offset to your policy firming. Without that offset, you may be concerned that the swing of both monetary and fiscal policies toward more restraint may slow the growth of spending next year even more than is built into the Greenbook.*

FOMC Forward Guidance: 2011 - 2012

August 9, 2011 FOMC Meeting: Negative anticipated monetary policy shock

Statement. *"The Committee currently anticipates that economic conditions–including low rates of resource utilization and a subdued outlook for inflation over the medium run–are likely to warrant exceptionally low levels for the federal funds rate at least through mid-2013. [. . .] The Committee discussed the range of policy tools available to promote a stronger economic recovery in a context of price stability. It will continue to assess the economic outlook in light of incoming information and is prepared to employ these tools as appropriate."*

Minutes. *"In the discussion of monetary policy for the period ahead, most members agreed that the economic outlook had deteriorated by enough to warrant a Committee response at this meeting. While all felt that monetary policy could not completely address the various strains on the*

economy, most members thought that it could contribute importantly to better outcomes in terms of the Committee's dual mandate of maximum employment and price stability. In particular, some members expressed the view that additional accommodation was warranted because they expected the unemployment rate to remain well above, and inflation to be at or below, levels consistent with the Committee's mandate."

From the transcripts it is clear that the statement chosen, over the various alternative considered reflects the intention of providing additional stimulus to the economy. This for instance transpires by the Secretary of the FOMC, Mr. English, description of the rationale for adopting the statement that was eventually released. Specifically, he mentions *"With the incoming data over the intermeeting period suggesting an even weaker outlook for economic activity, some participants may feel that it is now appropriate to provide additional accommodation, as in alternative A. [...] Regarding monetary policy, alternative A would include two steps to provide additional accommodation. First, paragraph 3 would provide more-explicit forward guidance about the expected path for the federal funds rate by specifying that exceptionally low levels were likely "at least through mid2013." Second, [...] Although expectations may have moved in recent days, market participants would probably still be surprised by the adoption of alternative A. Interest rates and the foreign exchange value of the dollar would likely fall, and stock prices would probably increase."*

January 24-25, 2012 FOMC Meeting: Negative anticipated monetary policy shock

Statement. *"To support a stronger economic recovery and to help ensure that inflation, over time, is at levels consistent with the dual mandate, the Committee expects to maintain a highly accommodative stance for monetary policy. In particular, the Committee decided today to keep the target range for the federal funds rate at 0 to 1/4 percent and currently anticipates that economic conditions–including low rates of resource utilization and a subdued outlook for inflation over the medium run–are likely to warrant exceptionally low levels for the federal funds rate at least through late 2014."*

Minutes. *"With respect to the statement to be released following the meeting, members agreed that only relatively small modifications to the first two paragraphs were needed to reflect the incoming information and the modest changes to the economic outlook implied by the recent data. In light of the economic outlook, almost all members agreed to indicate that the Committee expects to maintain a highly accommodative stance for monetary policy and currently anticipates that economic*

conditions–including low rates of resource utilization and a subdued outlook for inflation over the medium run–are likely to warrant exceptionally low levels for the federal funds rate at least through late 2014, longer than had been indicated in recent FOMC statements. In particular, several members said they anticipated that unemployment would still be well above their estimates of its longer-term normal rate, and inflation would be at or below the Committee's longer-run objective, in late 2014. It was noted that extending the horizon of the Committee's forward guidance would help provide more accommodative financial conditions by shifting downward investors' expectations regarding the future path of the target federal funds rate. Some members underscored the conditional nature of the Committee's forward guidance and noted that it would be subject to revision in response to significant changes in the economic outlook."

From the transcripts it is clear that the statement chosen, over the various alternative considered reflects the intention of providing additional stimulus to the economy. This for instance transpires by the Secretary of the FOMC, Mr. English, description of the rationale for adopting the statement that was eventually released. Specifically, he mentions *"However, your SEP submissions suggest that you see only modest economic growth over coming quarters, with the unemployment rate declining even more slowly than you anticipated at the time of the November SEP. [...] With regard to inflation, the SEP results suggest that most of you see inflation as likely to be subdued. Against this backdrop, you may wish to ease financial conditions a bit further by providing forward guidance that pushes out the expected date of the liftoff of the federal funds rate. You may also think it helpful to provide clear guidance regarding the economic conditions that would warrant retaining the current very low level of the federal funds rate. [...] Market participants see significant odds that the Committee will raise the funds rate by the middle of 2014, and so specifying a date of "at least through late 2014" would likely push out somewhat expectations for the date of the first increase in the funds rate. Longer-term interest rates would likely decline a little, equity prices could rise some, and the foreign exchange value of the dollar might soften. "*

September 12-13, 2012 FOMC Meeting: Negative anticipated monetary policy shock

Statement. *" The Committee is concerned that, without further policy accommodation, economic growth might not be strong enough to generate sustained improvement in labor market conditions. Furthermore, strains in global financial markets continue to pose significant downside risks to the economic outlook. The Committee also anticipates that inflation over the medium term likely would*

run at or below its 2 percent objective.[...] To support continued progress toward maximum employment and price stability, the Committee expects that a highly accommodative stance of monetary policy will remain appropriate for a considerable time after the economic recovery strengthens. In particular, the Committee also decided today to keep the target range for the federal funds rate at 0 to 1/4 percent and currently anticipates that exceptionally low levels for the federal funds rate are likely to be warranted at least through mid-2015."

Minutes. *"Members generally judged that without additional policy accommodation, economic growth might not be strong enough to generate sustained improvement in labor market conditions. [...] Inflation had been subdued, even though the prices of some key commodities had increased recently. Members generally continued to anticipate that, with longer-term inflation expectations stable and given the existing slack in resource utilization, inflation over the medium term would run at or below the Committee's longer-run objective of 2 percent.*

In their discussion of monetary policy for the period ahead, members generally expressed concerns about the slow pace of improvement in labor market conditions and all members but one agreed that the outlook for economic activity and inflation called for additional monetary accommodation. Members agreed that such accommodation should be provided through both a strengthening of the forward guidance regarding the federal funds rate and purchases of additional agency MBS at a pace of $40 billion per month. [...] With regard to the forward guidance, the Committee agreed on an extension through mid-2015, in conjunction with language in the statement indicating that it expects that a highly accommodative stance of policy will remain appropriate for a considerable time after the economic recovery strengthens. That new language was meant to clarify that the maintenance of a very low federal funds rate over that period did not reflect an expectation that the economy would remain weak, but rather reflected the Committee's intention to support a stronger economic recovery."

The presentation material for this meeting associated with the alternative monetary policy options, presents simulations which highlight how the adopted statement was meant to provide additional stimulus to the economy. In fact, the secretary of the FOMC, Mr. English, states *"The black solid lines in the charts depict the experimental consensus forecast, which is conditioned on unchanged policy. The red dashed lines show the results of a more accommodative policy, one that's consistent with an alternative B that shifts back the date in the forward guidance to mid-2015 and, in addition to $30 billion of MBS purchases a month and completion of the MEP this year,*

involves buying longer-term securities at a rate of $75 billion a month through the middle of next year. [...]] So long as the public correctly anticipates that the Committee will follow this policy, the result is a more rapid economic recovery that takes the unemployment rate to 6.3 percent by the end of 2015, about 1/2 percentage point lower than in the consensus baseline; the inflation rate (shown at the bottom right) is a bit higher but remains near your 2 percent longer-run objective. [...] If, in reviewing the outlook under the unchanged policy, the Committee views the likely outcomes for employment and inflation as inconsistent with its mandate, it might choose to ease policy by strengthening the forward guidance in the statement and engaging in additional asset purchases. Alternative B, on page 5, may offer the Committee an attractive approach to moving strongly in that direction at this meeting, but without committing now to a large, discrete purchase program. "

J Additional Figures

Figure J.1 analyzes the consequences of the alternative treatments of uncertainty for the density forecasts for the structural scenario: "What is the likely path of output and inflation, if monetary policy shocks lower the fed funds rate to 1% and keep it there for two years?"

Figure J.2 shows an alternative scenario for the exercise of returning to the price level target using only unanticipated shocks. The results indicate that a more volatile interest rate path would be required to achieve the same objective. The fed funds rate would be cut to 1% by mid-2020 and then increased to 4.6% in 2023. During that year the probability of negative annual GDP growth would exceed 50%. Figure J.3 instead shows a scenario in which both anticipated and unanticipated shocks are used, but the price level returns to the target in 7 years instead of 4 years.

Figure J.1: Comparison under Different Treatments of Uncertainty

(a) No Uncertainty ($\mathbf{\Omega}_f = \mathbf{0}, \mathbf{\Omega}_g = \mathbf{0}$)

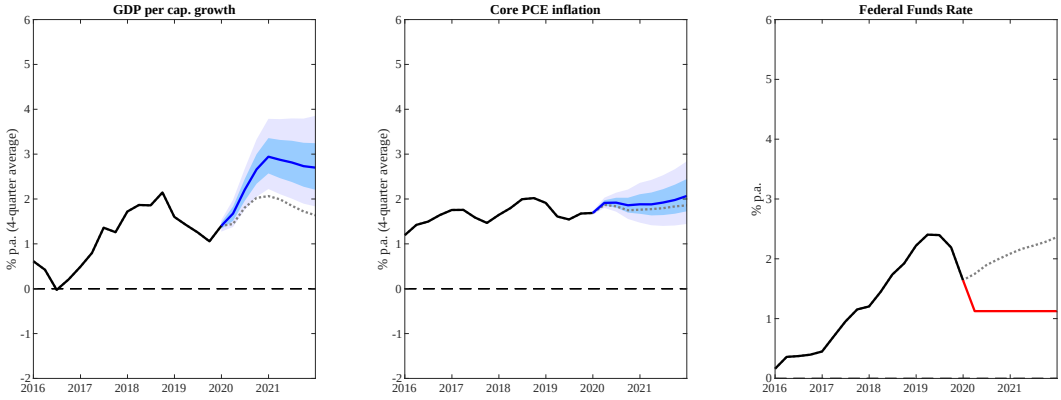

(b) Only Uncertainty about the Non-Driving Shocks ($\mathbf{\Omega}_f = \mathbf{0}, \mathbf{\Omega}_g = \mathbf{I}$)

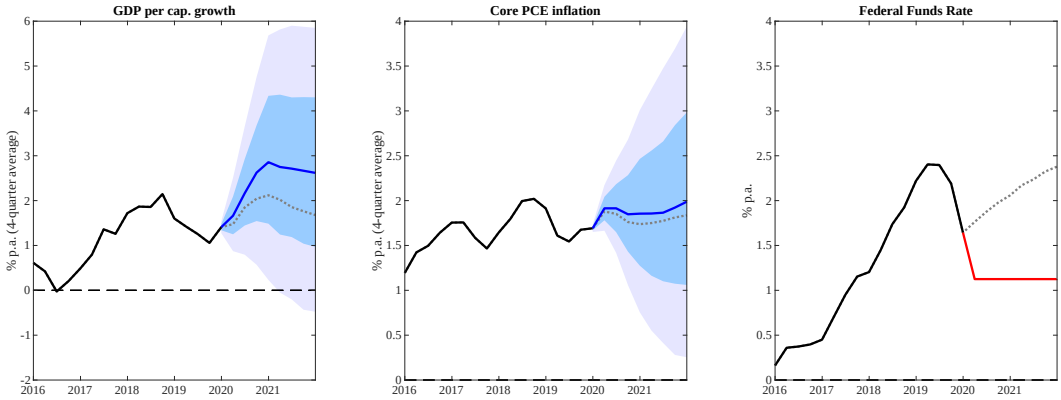

(c) Full Uncertainty ($\mathbf{\Omega}_f = \mathbf{D}\mathbf{D}', \mathbf{\Omega}_g = \mathbf{I}$)

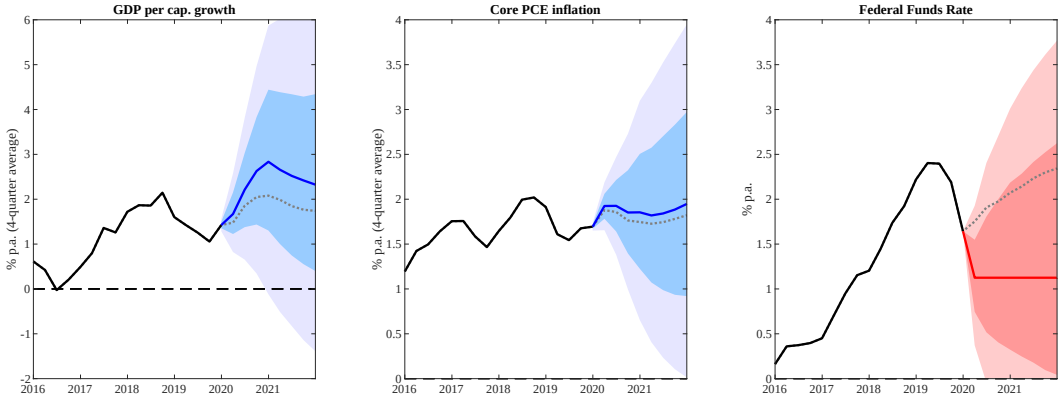

Note: For each panel, the solid black lines represent actual data, the solid red line is the conditioning assumption on the fed funds rate, the solid blue line is the median forecast for the remaining variables and periods, and the blue shaded areas denote the 40 and 68% pointwise credible sets around the forecasts. The dotted gray lines represent the median unconditional forecast.

Figure J.2: Returning to the Price Level Target: Unanticipated Shocks Only

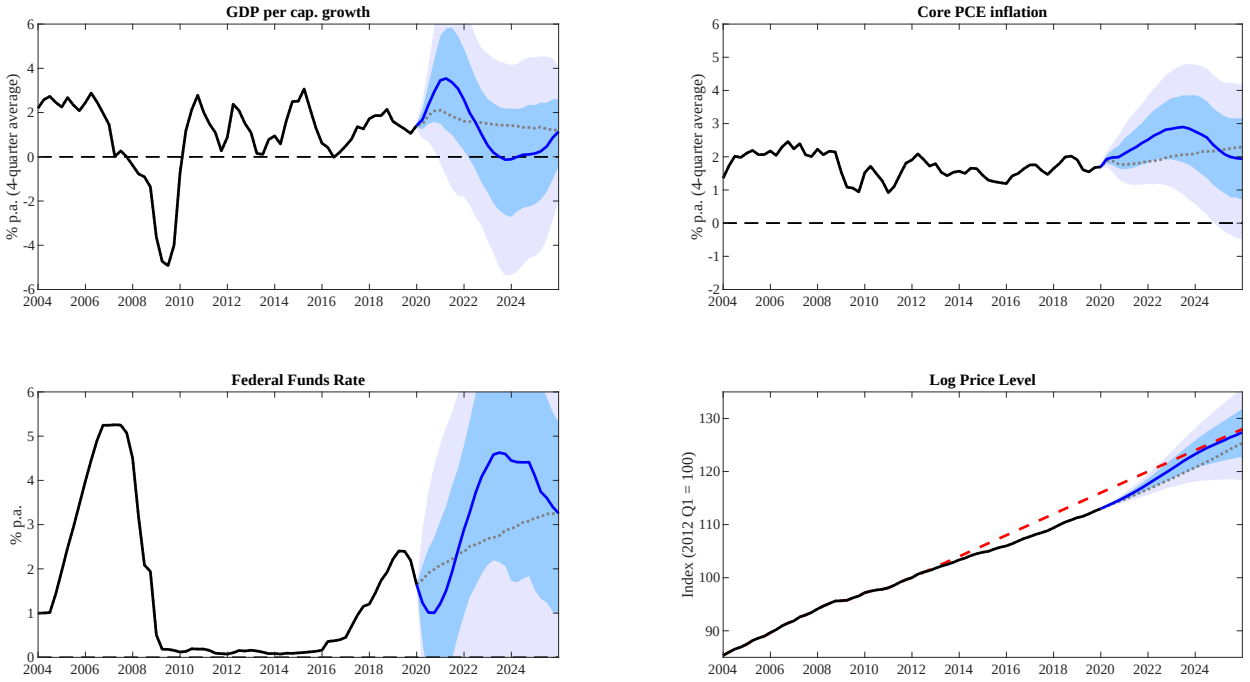

Note: For each panel, the solid black lines represent actual data. The solid blue lines and associated shaded areas are the median, 40 and 68% pointwise credible sets around the forecasts of the other variables given such that the constraint that inflation returns to the price level target within a five-year horizon is satisfied. The dotted gray lines represent the median unconditional forecast.

Figure J.3: Returning to the Price Level Target: Seven Year Horizon

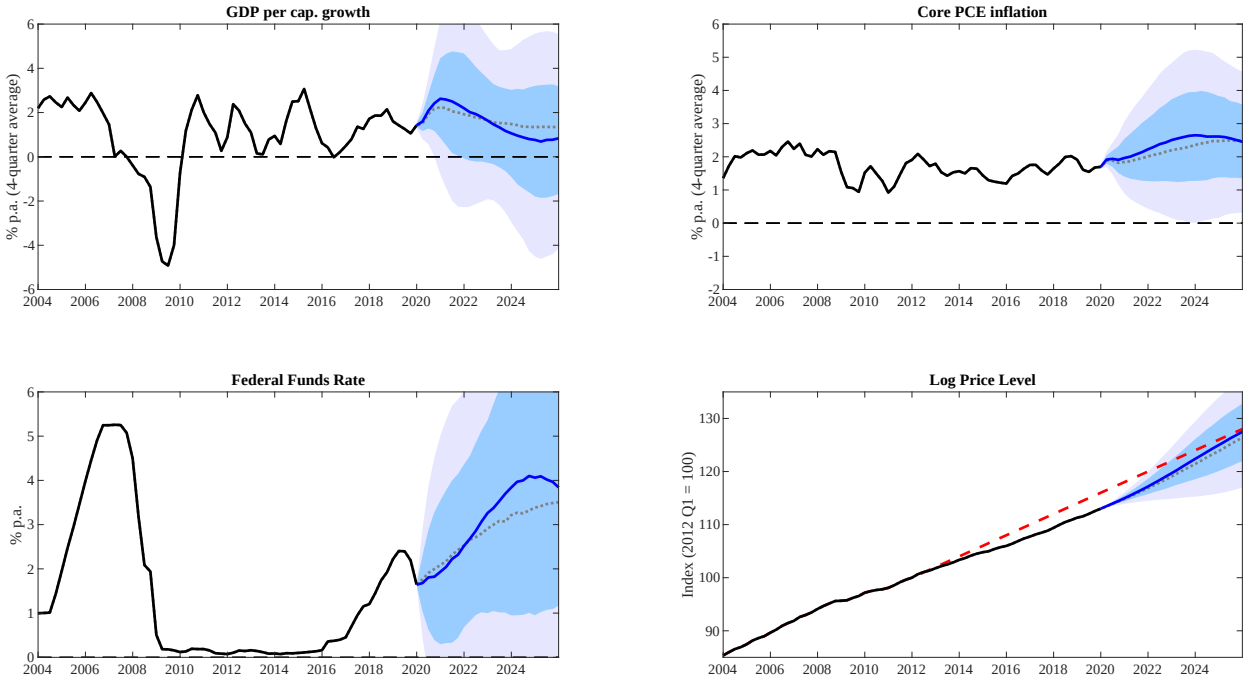

Note: For each panel, the solid black lines represent actual data. The solid blue lines and associated shaded areas are the median, 40 and 68% pointwise credible sets around the forecasts of the other variables given such that the constraint that inflation returns to the price level target within a seven-year horizon is satisfied. The dotted gray lines represent the median unconditional forecast.

References

Andersson, M. K., S. Palmqvist, and D. F. Waggoner (2010): "Density Conditional Forecasts in Dynamic Multivariate Models," Sveriges Riksbank Working Paper Series 247, Sveriges Riksbank.

Antolin-Diaz, J. and J. F. Rubio-Ramirez (2018): "Narrative Sign Restrictions for SVARs," *American Economic Review*, 108, 2802–29.

Arias, J. E., J. F. Rubio-Ramirez, and D. F. Waggoner (2018): "Inference Based on SVARs Identified with Sign and Zero Restrictions: Theory and Applications," *Econometrica*, 86, 685–720.

Banbura, M., D. Giannone, and M. Lenza (2015): "Conditional forecasts and scenario analysis with vector autoregressions for large cross-sections," *International Journal of Forecasting*, 31, 739–756.

Baumeister, C. and J. D. Hamilton (2015): "Sign restrictions, structural vector autoregressions, and useful prior information," *Econometrica*, 83, 1963–1999.

Ben-israel, A. and T. N. E. Greville (2001): *Generalized Inverses: Theory and Applications*, PublisherSpringer-Verlag New York, second ed.

Brockwell, P., R. Davis, and H. Salehi (1990): "A state-space approach to transfer-function modeling," *Statistical Inference in Stochastic Processes*, 6, 233.

Brockwell, P. J. and R. A. Davis (1991): "Time Series: Theory and Methods," *New York: Springer-Verlag.*, 66, 119–128.

Camacho, M. and G. Perez-Quiros (2010): "Introducing the euro-sting: Short-term indicator of euro area growth," *Journal of Applied Econometrics*, 25, 663–694.

Campbell, J. R., C. L. Evans, J. D. Fisher, and A. Justiniano (2012): "Macroeconomic Effects of Federal Reserve Forward Guidance," *Brookings Papers on Economic Activity*, 43, 1–80.

Canova, F. and G. D. Nicolo (2002): "Monetary disturbances matter for business fluctuations in the G-7," *Journal of Monetary Economics*, 49, 1131–1159.

Clarida, R. H. and D. Coyle (1984): "Conditional Projection by Means of Kalman Filtering," NBER Technical Working Papers 0036, National Bureau of Economic Research, Inc.

deJong, P. (1988): "A cross-validation filter for time series models." *Biometrika*, 75, 594–600.

Del Negro, M. (2011): "Bayesian Macroeconometrics," in *The Oxford Handbook of Bayesian Econometrics*.

Doan, T., R. B. Litterman, and C. A. Sims (1986): "Forecasting and conditional projection using realistic prior distribution," Staff Report 93, Federal Reserve Bank of Minneapolis.

Durbin, J. and S. J. Koopman (2012): *Time Series Analysis by State Space Methods: Second Edition*, Oxford University Press.

Faust, J. (1998): "The robustness of identified VAR conclusions about money," International Finance Discussion Papers 610, Board of Governors of the Federal Reserve System (U.S.).

Fernández-Villaverde, J., J. F. Rubio-Ramírez, T. J. Sargent, and M. W. Watson (2007): "ABCs (and Ds) of Understanding VARs," *American Economic Review*, 97, 1021–1026.

Golub, G. H. and C. F. Van Loan (1996): *Matrix Computations*, The Johns Hopkins University Press, third ed.

Karlsson, S. (2013): "Forecasting with Bayesian vector autoregression," in *Handbook of economic forecasting*, Elsevier, vol. 2, 791–897.

Kilian, L. and H. Lutkepohl (2017): *Structural Vector Autoregressive Analysis*, Cambridge University Press.

Koch, K. (2007): "Gibbs sampler by sampling-importance-resampling," *Journal of Geodesy*, 81, 581–591.

Mariano, R. S. and Y. Murasawa (2003): "A new coincident index of business cycles based on monthly and quarterly series," *Journal of Applied Econometrics*, 18, 427–443.

Penrose, R. (1956): "On best approximate solutions of linear matrix equations," *Mathematical Proceedings of the Cambridge Philosophical Society*, 52, 17–19.

Rubio-Ramirez, J. F., D. F. Waggoner, and T. Zha (2010): "Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference," *Review of Economic Studies*, 77, 665–696.

Searle, S. R. (1982): *Matrix Algebra Useful for Statistics*, John Wiley & Sons.

Sims, C. A. (1993): "A Nine-Variable Probabilistic Macroeconomic Forecasting Model," in *Business Cycles, Indicators and Forecasting*, National Bureau of Economic Research, Inc, NBER Chapters, 179–212.

Smets, F. and R. Wouters (2007): "Shocks and frictions in US business cycles: A Bayesian DSGE approach," *American economic review*, 97, 586–606.

Uhlig, H. (2005): "What are the effects of monetary policy on output? Results from an agnostic identification procedure," *Journal of Monetary Economics*, 52, 381–419.

Waggoner, D. F. and T. Zha (1999): "Conditional Forecasts in Dynamic Multivariate Models," *Review of Economics and Statistics*, 81, 639–651.