Structural Scenario Analysis with SVARs*

Juan Antolín-Díaz — Fulcrum Asset Management

Ivan Petrella — University of Warwick

Juan F. Rubio-Ramírez†

Emory University

Federal Reserve Bank of Atlanta

Abstract

In the context of vector autoregressions, conditional forecasts are typically constructed by specifying the future path of one or more variables while remaining silent about the structural shocks that might have caused the path. However, in many cases, researchers may be interested in identifying a structural vector autoregression and choosing which structural shock is driving the path of the conditioning variables. This would allow researchers to create a "structural scenario" that can be given an economic interpretation. In this paper we show how to construct structural scenarios and develop efficient algorithms to implement our methods. We show how structural scenario analysis can lead to results that are very different from, but complementary to, those of the traditional conditional forecasting exercises. We also propose an approach to assess and compare the plausibility of alternative scenarios. We illustrate our methods by applying them to two examples: comparing alternative monetary policy options and stress testing the reaction of bank profitability to an economic recession.

Keywords: Conditional forecasts, Bayesian methods, probability distribution, SVARs

JEL Classification Numbers: C32, C53, E47.

*We are grateful to Gavyn Davies and Jonas Hallgren for helpful comments and suggestions. Yad Selvakumar provided excellent research assistance.

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

1 Introduction

A key question in applied macroeconomics and policy analysis is: "If, for the next few quarters, variable x follows alternative paths, how do the forecasts of other variables change?" These alternative forecasts are called conditional forecasts. Common uses of conditional forecasts include assessing the path of macroeconomic variables to alternative scenarios for a monetary policy instrument, incorporating external information such as data from futures prices to condition on the path of oil prices or other financial variables, and "stress testing," e.g. assessing the reaction of asset prices or bank profits to an economic recession. Waggoner and Zha (1999) provide methods for computing conditional density forecasts in the context of vector autoregression (VAR) models, Banbura, Giannone, and Lenza (2015) extend the application to large systems, and Andersson, Palmqvist, and Waggoner (2010) extend the analysis to the case when there is uncertainty about the paths of the conditioning variables.

As an illustrative example, consider a three-variable VAR with output growth, inflation, and the policy interest rate that will be analyzed in Section 5.1. A policy maker might want to use this VAR to ask the question: "What is the likely path of output and inflation, given that the fed funds rate is kept at zero for two years?" We will call this exercise a "conditional-on-observables forecast." It can be computed using the methods of Waggoner and Zha (1999) and it is presented in Panel (a) in Figure 1. The dotted lines represent the median and 68 percent credible intervals around the unconditional forecasts and the blue line and shaded areas represent the equivalent quantities for the conditional forecast. As can be seen from the figure, the main result of conditioning the forecast on zero interest rates for two years is that inflation is about half a percent lower than in the unconditional forecast.

At first this result might appear puzzling, as easy monetary policy is usually thought to stimulate output and inflation. However, one should keep in mind that because a large proportion of the movements in the federal funds rate represents the systematic reaction of

Figure 1: Answers to Two Alternative Questions

(a) Conditional forecasting as described in Waggoner and Zha (1999)

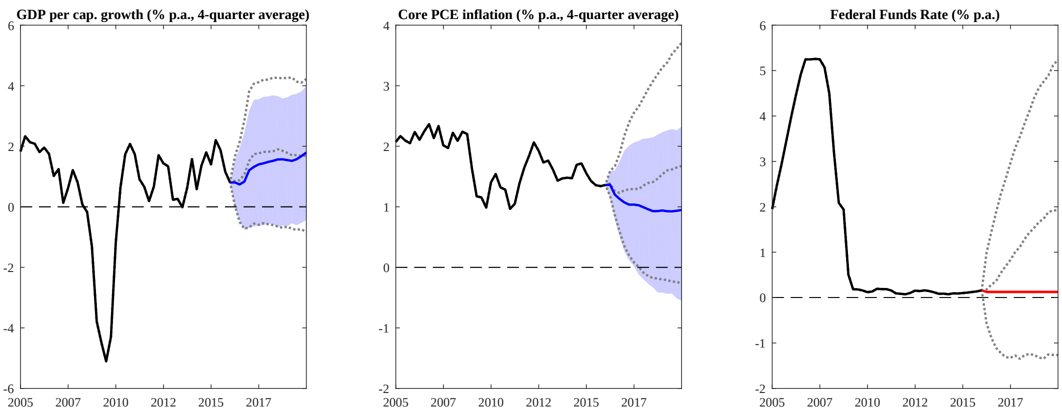

(b) Structural Scenario Analysis

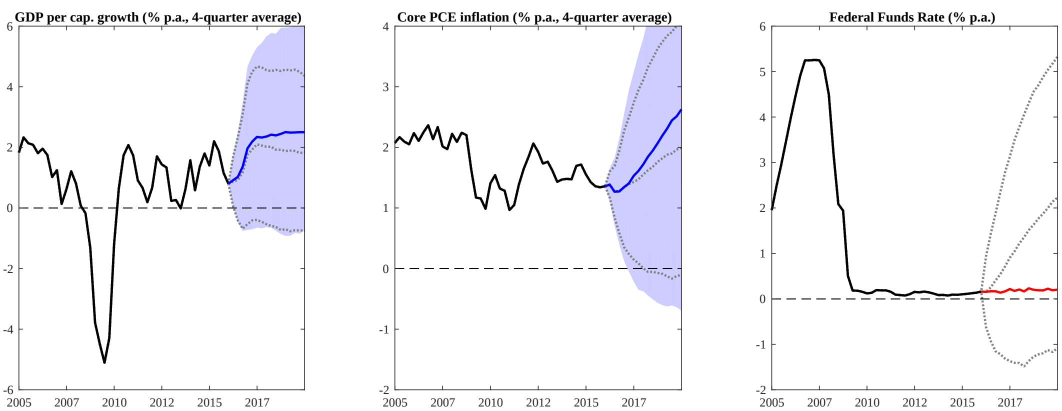

Note: For each column, the solid black lines represent actual data, the solid red line is the conditioning assumption on the observables, the solid blue line is the median forecast for the remaining variables and periods, and the blue shaded areas denote the 68 percent pointwise credible sets around the forecasts. The dotted black lines represent the median and contours of the 68 percent credible set around the unconditional forecast.

the Fed to output and inflation developments (see, e.g., Leeper, Sims, and Zha (1996)), the unconditional correlation between the interest rate and both inflation and output is strong and positive in the data. In other words, the federal funds rates are usually low because the Fed is systematically responding to low output and inflation. The result is not so puzzling anymore once one takes into account this correlation and the fact that the conditional forecast, as described in Waggoner and Zha (1999), is really answering the question: "What is the most likely set of circumstances under which the Fed might keep the federal funds rate at zero for two years?"[1]

In this paper, we develop tools to answer a different question. Going back to the monetary policy example, we want a response to the alternative query: "What is the likely path of output and inflation, if a sequence of monetary policy shocks keeps the federal funds rate at zero for two years?" We call this exercise a "structural scenario analysis." The result is displayed in Panel (b) of Figure 1. Note that the conditioning path for the federal funds rate is identical to that of Panel (a), but output and inflation are now slightly higher than in the unconditional forecast.[2] The reason for the discrepancy is that, in the case of structural scenario analysis, the federal funds rate deviates from the unconditional forecast as a result of a series of monetary policy shocks unfolding over the forecast horizon. A key difference with conditional forecasting, as described in Waggoner and Zha (1999), is that the latter does not require identifying the structural shocks, while our structural scenario analysis does. In other words, the conditional-on-observables forecast can be done in a VAR context, while a structural scenario analysis exercise needs a structural vector autoregression (SVAR). As expected, the results will critically depend on the identifying restrictions used to identify the SVAR.

[1] Campbell, Evans, Fisher, and Justiniano (2012) call this exercise "Delphic" forward guidance. It provides a forecast "of macroeconomic performance and likely or intended monetary policy actions based on the policymaker's potentially superior information about future macroeconomic fundamentals and its own policy goals."

[2] Similar to Campbell, Evans, Fisher, and Justiniano's (2012) "Odyssean" forward guidance, the exercise does not reveal any information about future aggregate demand and supply shocks, but rather can be interpreted as a commitment by the policy maker to keep the policy rate at zero whatever these shocks are.

We are not the first to propose conditioning forecasts on structural shocks. Baumeister and Kilian (2014) perturb an SVAR of the oil market with structural shocks to trace out the impact on oil prices. This practice is also used routinely in the DSGE literature, where a fully specified structural model is also available (see Del Negro and Schorfheide (2013)). Relative to these methods, our proposal carries several innovations. First, we show that fewer restrictions on the structural shocks are needed. While Baumeister and Kilian (2014) set the "non-restricted" structural shocks to zero, we only set their distribution to be standard normal. This will allow us to consider the uncertainty regarding the "non-restricted" structural shocks over the forecast horizon. Second, we extend the existing Bayesian methods to set and partially identified SVARs. This will allow us to consider both parameter and model uncertainty.[3] Of course, the same approach can be used for exactly and fully identified SVARs. Lastly, we propose a way of assessing how plausible (or implausible) a structural scenario is. This tool offers a simple way of ranking alternative structural scenarios and a simple metric to evaluate whether a particular scenario can be evaluated within a linear VAR setting (see Leeper and Zha (2003)).

We illustrate our technique with two examples. First, we further develop the monetary example above and explore the intuition behind the results in terms of the underlying structural shocks. We show that, for the exact same path for the conditioning variables, the scenarios can be very different depending on which structural shock is assumed to drive the scenario. Second, we consider a larger VAR 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. In particular, identifying only one structural shock, we contrast two alternative scenarios in which the same recession is generated by a financial shock or by non-financial shocks. In this setting we highlight how a recession of the same size but of different origin has very different effects on key financial indicators. This highlights the

[3] The existing procedures, such as Baumeister and Kilian (2014) and Clark and McCracken (2014), have tended to disregard parameter uncertainty. Waggoner and Zha (1999) find that ignoring parameter uncertainty can potentially result in misleading conditional forecasts.

importance of considering different structural interpretations of the same conditional scenario in stress-testing exercises.

The rest of this paper is organized as follows. Section 2 presents the general econometric framework. Section 3 formalizes the concept of structural scenario analysis, distinguishing it from conditional-on-observables and conditional-on-shocks forecast. This section also introduces a way to measure the plausibility of the different scenarios under consideration. Section 4 provides algorithms to implement our techniques. Section 5 illustrates our techniques using two applications: a small monetary SVAR for analyzing the effects of a monetary policy tightening and a larger SVAR that we use for analyzing the effects of a recession on asset prices and bank profitability. Section 6 offers some concluding remarks.

2 Econometric framework

Consider the structural vector autoregression (SVAR) with the general form

$$\mathbf{y}_t'\mathbf{A}_0 = \sum_{\ell=1}^{p} \mathbf{y}_{t-\ell}'\mathbf{A}_\ell + \mathbf{d} + \boldsymbol{\varepsilon}_t' \text{ for } 1 \leq t \leq T, \tag{1}$$

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 a $1 \times n$ vector of parameters, p is the lag length, and T is the sample size. The vector of structural shocks $\boldsymbol{\varepsilon}_t$, conditional on past information and the initial conditions $[\mathbf{y}_0, \ldots, \mathbf{y}_{1-p}]$, is Gaussian with mean zero and covariance matrix $\mathbf{I}_n$, the $n \times n$ identity matrix. The model described in Equation (1) can be written as

$$\mathbf{y}_t'\mathbf{A}_0 = \mathbf{x}_t'\mathbf{A}_+ + \boldsymbol{\varepsilon}_t' \text{ for } 1 \leq t \leq T, \tag{2}$$

where $\mathbf{A}_+' = \begin{bmatrix}\mathbf{A}_1' & \cdots & \mathbf{A}_p' & \mathbf{d}'\end{bmatrix}$ and $\mathbf{x}_t' = \begin{bmatrix}\mathbf{y}_{t-1}', \ldots, \mathbf{y}_{t-p}', 1\end{bmatrix}$ for $1 \leq t \leq T$. The dimension of $\mathbf{A}_+$ is $m \times n$ and the dimension of $\mathbf{x}_t$ is $m \times 1$, where $m = np + 1$. The reduced-form representation implied by Equation (2) is

$$\mathbf{y}_t' = \mathbf{x}_t'\mathbf{B} + \mathbf{u}_t' \text{ for } 1 \leq t \leq T, \tag{3}$$

where $\mathbf{B} = \mathbf{A}_+\mathbf{A}_0^{-1}$, $\mathbf{u}_t' = \boldsymbol{\varepsilon}_t'\mathbf{A}_0^{-1}$, and $\mathbb{E}\left[\mathbf{u}_t\mathbf{u}_t'\right] = \boldsymbol{\Sigma} = \left(\mathbf{A}_0\mathbf{A}_0'\right)^{-1}$. 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 $\boldsymbol{\varepsilon}_t'$ are the structural shocks. Note that the shocks are orthogonal and have an economic interpretation, while the innovations are, in general, correlated and do not have an interpretation.

Finally, the SVAR can alternatively be written in terms of the orthogonal reduced-form parameterization. This parameterization is particularly convenient for simulation, and is given by the following equation

$$\mathbf{y}_t' = \mathbf{x}_t'\mathbf{B} + \boldsymbol{\varepsilon}_t' h(\boldsymbol{\Sigma})\mathbf{Q}^{-1} \text{ for } 1 \leq t \leq T, \tag{4}$$

where $h(\boldsymbol{\Sigma})$ is any decomposition of the covariance matrix $\boldsymbol{\Sigma}$, such as the Cholesky decomposition, that satisfies $h(\boldsymbol{\Sigma})'h(\boldsymbol{\Sigma}) = \boldsymbol{\Sigma}$, and $\mathbf{Q}$ is an $n \times n$ orthogonal matrix. The orthogonal reduced-form parameterization makes clear how the structural parameters depend on the reduced-form parameters $\mathbf{B}$ and $\boldsymbol{\Sigma}$ together with an orthogonal rotation matrix $\mathbf{Q}$. For full details on the mapping between the structural and the orthogonal-reduced form parameterizations, we refer to Arias, Rubio-Ramirez, and Waggoner (2016b). Suffice to say here that given the reduced-form parameters and a decomposition h, one can consider each value of the orthogonal matrix $\mathbf{Q}$ as a particular choice of structural parameters. This implies that there is an identification problem. Each reduced-form parameter is linked to several structural parameters. It is well-known that to solve the identification problem, one often imposes restrictions on either the structural parameters or some function of the structural parameters, such as the impulse response functions.

2.1 Unconditional forecasting

Assume that we want to forecast the observables for h periods ahead using Equations (2)-(4). Conditional on the history of observables $\mathbf{y}^T = \left(\mathbf{y}'_{1-p} \ldots \mathbf{y}'_T\right)'$, the forecast of $\mathbf{y}'_{T+1,T+h} = \left(\mathbf{y}'_{T+1} \ldots \mathbf{y}'_{T+h}\right)$ can be rewritten as:

$$\mathbf{y}'_{T+1,T+h} = \mathbf{b}'_{T+1,T+h} + \boldsymbol{\varepsilon}'_{T+1,T+h}\mathbf{M} \text{ for all } 1 < t < T \text{ and all } h > 0. \tag{5}$$

The vector $\mathbf{b}_{t+1,t+h}$ and matrix $\mathbf{M}$ depend on the parameters of the model and their definitions are given in Appendix A. The first term, $\mathbf{b}_{T+1,T+h}$ is deterministic and gives the dynamic forecast in the absence of future structural shocks, whereas $\boldsymbol{\varepsilon}'_{T+1,T+h}\mathbf{M}$ is a stochastic part, reflecting the structural shocks unfolding over the forecast horizon, with $\mathbf{M}$ capturing the associated impulse response functions' (IRFs) coefficients. Given Equation (5) the unconditional forecast $\mathbf{y}'_{T+1,T+h}$ is distributed as

$$\mathbf{y}_{T+1,T+h} \sim \mathcal{N}\left(\mathbf{b}_{T+1,T+h}, \mathbf{M}'\mathbf{M}\right). \tag{6}$$

Importantly, in Appendix A we show that while $\mathbf{M}$ depends on the structural parameters, $\mathbf{M}'\mathbf{M}$ only depends on the reduced-form parameters, i.e., given the reduced-form parameters, the choice of $\mathbf{Q}$ is irrelevant for the unconditional forecast $\mathbf{y}'_{T+1,T+h}$. Thus, identification is not an issue when doing unconditional forecasting.

It is also important to notice that the distribution of $\boldsymbol{\varepsilon}_{T+1,T+h}$ compatible with Equation (6) is

$$\boldsymbol{\varepsilon}_{T+1,T+h} \sim \mathcal{N}\left(\mathbf{0}_{nh\times 1}, \mathbf{I}_{nh}\right), \tag{7}$$

where $\mathbf{0}_{nh\times 1}$ is the null vector of dimension $nh \times 1$ and $\mathbf{I}_{nh}$ is a conformable identity matrix. Hence, unconditional forecasting does not restrict the distribution of $\boldsymbol{\varepsilon}_{T+1,T+h}$.

3 Conditional forecasting and structural scenario

It is often the case that one wants to incorporate external information into a forecast, such as in the examples of conditional forecasts illustrated in the introduction. In this section we describe different ways of incorporating conditional assumptions into a forecast. In particular, we first consider the exercise conditioning on a path for a subset of the observables, second conditioning on a path for the structural shocks, finally we introduce the concept of structural scenario analysis, which combines restrictions on both observables and structural shocks. We will call the first exercise conditional-on-observables forecasting. The second exercise will be called conditional-on-shocks forecasting. The final exercise will be called structural scenario analysis. The conditional-on-observables forecasting exercise is what most people refer to as conditional forecasting. The conditional-on-shocks forecasting is just an intermediate step that we use to develop our structural scenario analysis. Despite their different interpretations, the three variants can all be written as linear restrictions on the path of future observables. Therefore, we first describe the general framework for incorporating these types of linear restrictions on the forecast and then we proceed to show how each of the three exercises can be written as a special case of the former.

3.1 General framework for conditional forecasts

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

$$\mathbf{C}\mathbf{y}_{T+1,T+h} \sim \mathcal{N}\left(\mathbf{f}_{T+1,T+h}, \mathbf{\Omega}_f\right), \tag{8}$$

where $\mathbf{C}$ is a pre-specified matrix of dimension $k \times nh$, with k denoting the number of restrictions (and $k \leq nh$) and the $1 \times k$ vector $\mathbf{f}'_{T+1,T+h}$ and the $k \times k$ matrix $\mathbf{\Omega}_f$ are the mean and variance restriction to the forecast of $\mathbf{y}'_{T+1,T+h}\mathbf{C}'$, respectively. This formulation accommodates density restrictions as well as the more common point restrictions in the special case of $\boldsymbol{\Omega}_f = \mathbf{0}_{nh}$.[4] Note that Equation (6) in turn implies that the unconditional forecast of $\mathbf{C}\mathbf{y}_{T+1,T+h}$ is normally distributed with mean $\mathbf{C}\mathbf{b}_{T+1,T+h}$ and variance $\mathbf{D}\mathbf{D}'$, where $\mathbf{D} = \mathbf{C}\mathbf{M}'$.

Andersson, Palmqvist, and Waggoner (2010) show that the distribution of the forecast $\mathbf{y}_{T+1,T+h}$, conditional on a set of linear restrictions given by Equation (8), can be written as

$$\mathbf{y}_{T+1,T+h} \sim \mathcal{N}\left(\boldsymbol{\mu}_y, \boldsymbol{\Sigma}_y\right), \tag{9}$$

with

$$\boldsymbol{\mu}_y = \mathbf{M}'\mathbf{D}^*\mathbf{f}_{T+1,T+h} + \mathbf{M}'\hat{\mathbf{D}}'\hat{\mathbf{D}}(\mathbf{M}')^{-1}\mathbf{b}_{T+1,T+h}, \tag{10}$$

and

$$\boldsymbol{\Sigma}_y = \mathbf{M}'\left[\mathbf{D}^*\boldsymbol{\Omega}_f(\mathbf{D}^*)' + \hat{\mathbf{D}}'\hat{\mathbf{D}}\right]\mathbf{M}, \tag{11}$$

where $\mathbf{D}^*$ is the generalized inverse of $\mathbf{D}$ and $\hat{\mathbf{D}}$ is any $(nh-k)\times nh$ such that its rows form an orthonormal basis for the null space of $\mathbf{D}$.[5] Going back to the general expression in Equation (5), it is then possible to show that in order to retrieve the distribution of the conditional forecasts in Equation (9), one is effectively adding to the deterministic component of the forecast, $\mathbf{b}_{T+1,T+h}$, structural shocks distributed as

$$\boldsymbol{\varepsilon}_{T+1,T+h} \sim \mathcal{N}\left(\boldsymbol{\mu}_\varepsilon, \boldsymbol{\Sigma}_\varepsilon\right), \tag{12}$$

[4] Here we focus on hard conditions on the distribution of the forecast and the case of soft conditioning is not explicitly considered. See Waggoner and Zha (1999) for a formal definition of the two different conditioning assumptions. Andersson, Palmqvist, and Waggoner (2010) show how the solution for the hard conditioning can be easily amended in order to deal with soft condition restrictions on the forecasts using truncated normals.

[5] Banbura, Giannone, and Lenza (2015) propose an alternative formulation of Equations 9-11 that uses the Kalman filter and smoother to recursively compute the conditional forecasts. These formulas lead to identical results but, as shown in Appendix C, the formulas above are computationally more efficient as long as the maximum forecast horizon, h, is not very large.

where

$$\boldsymbol{\mu}_{\varepsilon} = \mathbf{D}^{*}\mathbf{f}_{T+1,T+h} - \mathbf{D}^{*}\mathbf{C}\mathbf{b}_{T+1,T+h} \tag{13}$$

and

$$\boldsymbol{\Sigma}_{\varepsilon} = \mathbf{D}^{*}\boldsymbol{\Omega}_{f}(\mathbf{D}^{*})' + \hat{\mathbf{D}}'\hat{\mathbf{D}}. \tag{14}$$

Therefore conditional forecasts are associated with a distribution of the structural shocks over the forecast horizon that deviates from their unconditional distribution.[6] We now show how different methods to construct conditional forecasts are special cases of the general framework above.

3.2 Conditional-on-observables forecasting

The classical conditional forecasting exercises, such as those first introduced by Doan, Litterman, and Sims (1986), focus on calculating the forecast of observables conditional on an exogenously imposed path for a subset of observables. We will call this environment a conditional-on-observables forecast.

Let $\overline{\mathbf{C}}$ be a $k_o \times nh$ selection matrix formed by ones and zeros, with k_o denoting the number of restrictions (and $k_o \leq nh$). Conditional-on-observables restrictions are written as $\overline{\mathbf{C}}\mathbf{y}_{T+1,T+h} = \overline{\mathbf{f}}_{T+1,T+h}$ (see Waggoner and Zha, 1999) or more generally as density restrictions $\overline{\mathbf{C}}\mathbf{y}_{T+1,T+h} \sim \mathcal{N}\left(\overline{\mathbf{f}}_{T+1,T+h}, \overline{\boldsymbol{\Omega}}_{f}\right)$ (see Andersson, Palmqvist, and Waggoner, 2010). These types of restrictions are trivially expressed in terms of Equation (8), by making $\mathbf{C} = \overline{\mathbf{C}}$, $\mathbf{f}_{T+1,T+h} = \overline{\mathbf{f}}_{T+1,T+h}$ and $\boldsymbol{\Omega}_{f} = \overline{\boldsymbol{\Omega}}_{f}$. The solution for the conditional forecast of the entire set of observables is then given by Equation (9).

Note that, since the selection matrix $\overline{\mathbf{C}}$ does not depend on the parameters of the model, the distribution of the conditional-on-observables forecasting only depends on the reduced-form parameters. Hence, given a set of reduced-form parameters, the choice of structural

[6] See Appendix B in Andersson, Palmqvist, and Waggoner (2010) for details on Equations (9) and (12) and for a proof of the equivalence of this to the original solution of Waggoner and Zha (1999), under the restriction of $\boldsymbol{\Omega}_{f} = \mathbf{0}$.

parameters, i.e., $\mathbf{Q}$, is irrelevant for retrieving the distribution of the conditional forecast (see Waggoner and Zha, 1999, Proposition 1). Thus, identification is irrelevant when doing conditional-on-observables forecasting.

As shown before, it is the case that the conditional forecast is achieved by restricting the distribution of all structural shocks over the forecast horizon. Before the restrictions are imposed, we have that $\varepsilon_{T+1,T+h} \sim \mathcal{N}\left(\mathbf{0}_{nh\times 1}, \mathbf{I}_{nh}\right)$, while after the restrictions $\varepsilon_{T+1,T+h}$ are distributed as described in Equation (12). As we will see later, this implies that even if identification is not necessary to compute conditional-on-observable forecasts, one could back out the values of the structural shocks implied by the conditional forecast if an identification scheme was used.

3.3 Conditional-on-shocks forecasting

In this section we aim to construct forecasts for all the observables of the system conditioning on a particular path of the structural shocks over the forecast horizon. We will call this environment a conditional-on-shocks forecast.[7] For instance, 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 commonly used to produce conditional forecasts with DSGE models (see Del Negro and Schorfheide, 2013).

Formally, let Ξ be a $k_s \times nh$ selection matrix formed by ones and zeros, with k_s denoting the number of restrictions and h the number of periods ahead (and $k_s \leq nh$). Restrictions on the structural shocks can generally be written as

$$\Xi \boldsymbol{\varepsilon}_{T+1,T+h} \sim \mathcal{N}\left(\mathbf{g}_{T+1,T+h}, \boldsymbol{\Omega}_g\right), \tag{15}$$

where the $1 \times k_s$ vector $\mathbf{g}'_{T+1,T+h}$ and the conformable matrix $\boldsymbol{\Omega}_g$ denote the pre-specified

[7] Baumeister and Kilian (2014) call this possibility scenario analysis, while we call it conditional-on-shocks forecasting to make the comparison with Section 3.2 easier.

mean restriction to the forecast of $\boldsymbol{\varepsilon}'_{T+1,T+h}\boldsymbol{\Xi}'$ and the associated variance.[8] The structural shocks can always be retrieved from the observed variables, conditional on the structural parameters of the model. Specifically, Equation (5) implies that

$$\boldsymbol{\varepsilon}_{t+1,t+h} = (\mathbf{M}')^{-1}\mathbf{y}_{t+1,t+h} - (\mathbf{M}')^{-1}\mathbf{b}_{t+1,t+h}.$$

Therefore, the restriction on the structural shocks can be written as linear restrictions on the observables; specifically Equation (15) implies the following distribution on $\underline{\mathbf{C}}\mathbf{y}_{T+1,T+h}$

$$\underline{\mathbf{C}}\mathbf{y}_{T+1,T+h} \sim \mathcal{N}\left(\underline{\mathbf{C}}\mathbf{b}_{T+1,T+h} + \mathbf{g}_{T+1,T+h}, \boldsymbol{\Omega}_g\right), \tag{16}$$

where $\underline{\mathbf{C}} = \boldsymbol{\Xi}(\mathbf{M}')^{-1}$. Thus, the restricted forecast associated with Equation (15) is given by Equation (9) (with $\mathbf{C} = \underline{\mathbf{C}}$, $\mathbf{f}_{T+1,T+h} = \underline{\mathbf{C}}\mathbf{b}_{T+1,T+h} + \mathbf{g}_{T+1,T+h}$ and $\boldsymbol{\Omega}_f = \boldsymbol{\Omega}_g$). The crucial difference with respect to the conditional-on-variables forecasts is that the linear restrictions, $\underline{\mathbf{C}}$, now depend on the impulse response functions coefficients associated with the future shocks. Since $\mathbf{M}$ depends on the structural parameters, so will $\underline{\mathbf{C}}$, which implies that the identification will affect the conditional-on-shocks forecast of $\mathbf{y}_{T+1,T+h}$. The intuition is clear; in order to impose restrictions upon their path, structural shocks need to be identified. For this reason, and unlike the unconditional and conditional-on-variables forecasts, the conditional-on-shocks forecasting will depend on the structural parameters. Hence, given a set of reduced-form parameters, the choice of structural parameters, i.e., $\mathbf{Q}$, is now relevant. Moreover, it is worth noticing that, contrary to the conditional-on-observables forecasting, when conditional-on-shocks forecasting is used, the structural shocks that are not part of the conditioning exercise retain their unconditional (standard normal) distribution.[9]

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

[9] See Appendix B for the full derivation.

3.4 Structural scenario analysis

Conditional-on-shocks forecasting has the disadvantage that, since the structural shocks are unobserved, it is difficult to elicit a priori conditions on their values. In practice, the papers that use that method calibrate the value of the structural shocks to generate the desired impact on a particular variable, or iterate between the structural shocks and the observables until achieving that result (see, e.g., Baumeister and Kilian 2014, and Clark and McCracken 2014). These iterative procedures do not take into account the uncertainty associated with the conditional forecast. Here we show how the results of Sections 3.2 and 3.3 can be combined to approach this problem in a single step, which we call "structural scenario analysis."

A structural scenario is defined by the combination of a restriction on the path for one or more of the observables, together with a restriction that only a subset of the structural shocks can deviate from their unconditional distribution. The reader should remember that the conditional-on-observables method implied restrictions on all structural shocks. Here, only the structural shocks that are assumed to be drivers of the structural scenario are allowed to deviate from the standard normal distribution, whereas the rest of the structural shocks that are not explicitly part of the structural scenario are explicitly restricted to retain their unconditional distribution.

Let $\overline{\mathbf{C}}$ be a $k_o \times nh$ selection matrix formed by ones and zeros, with k_o denoting the number of restrictions on the observables. Let $\boldsymbol{\Xi}$ be a $k_s \times nh$ selection matrix formed by ones and zeros that selects the k_s structural shocks that are assumed not to be the key driver of the structural scenario, and therefore whose distribution is going to be restricted to be the same as their unconditional one.[10] Using the notation of Section 3.2, the restriction on the observables is implemented by imposing that

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

[10] It is required that $k_o + k_s \leq nh$. This implies that if we want to restrict m observables for the entire forecast horizon, we can keep unrestricted less than $n - m$ structural shocks.

While using the notation of Section 3.3, the restriction on the structural shocks is implemented by imposing that $\Xi\boldsymbol{\varepsilon}_{T+1,T+h}$ is distributed as follows

$$\Xi\boldsymbol{\varepsilon}_{T+1,T+h} \sim \mathcal{N}\left(\mathbf{0}_{k_s\times 1}, \mathbf{I}_{k_s}\right).$$

The latter, in turn, implies $\underline{\mathbf{C}}\mathbf{y}_{T+1,T+h} \sim \mathcal{N}\left(\underline{\mathbf{C}}\mathbf{b}_{T+1,T+h}, \mathbf{I}_{k_s}\right)$, where $\underline{\mathbf{C}} = \Xi(\mathbf{M}')^{-1}$. Taking the two sets of restrictions together, a structural scenario is a conditional forecast subject to the following restriction on the distribution of the observables over the forecast horizon

$$\mathbf{C}\mathbf{y}_{T+1,T+h} \sim \mathcal{N}\left(\begin{bmatrix} \mathbf{f}_{T+1,T+h} \\ \underline{\mathbf{C}}\mathbf{b}_{T+1,T+h} \end{bmatrix}, \begin{bmatrix} \mathbf{\Omega}_f & \mathbf{0}_{k_o\times k_s} \\ \mathbf{0}_{k_s\times k_o} & \mathbf{I}_{k_s} \end{bmatrix}\right), \tag{17}$$

with $\mathbf{C}' = [\overline{\mathbf{C}}', (\mathbf{M})^{-1}\Xi']$.

Given the restrictions in Equation (17), the distribution of the forecast $\mathbf{y}_{T+1,T+h}$ is defined by Equation (9) and it is associated with a conditional distribution of the shocks given by Equation (12). Observe that Equation (17) 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 shocks that do not form part of the structural scenario must retain their unconditional distribution. For the same reasons explained above, the structural scenario depends on the identification.

In Section 4 we describe algorithms to implement structural scenario analysis (and therefore conditional-on-observables and conditional-on-shocks forecasting as special cases) for set identified models. When describing those algorithms, we will make clear how the structural parameters and, therefore, the identification play a role in the forecast.

3.5 How plausible is the structural scenario?

When analyzing a structural scenario, it might be of interest to quantify its plausibility. In the previous subsection we have highlighted that the distribution of a conditional forecast is associated with a distribution of the structural shocks over the forecasting horizon that deviates from the unconditional distribution. Therefore, a structural scenario that requires a very unlikely distribution of structural shocks should be deemed implausible. This point was forcefully made by Leeper and Zha (2003).

We quantify how implausible a structural scenario is by determining how "far" the distribution of the structural shocks compatible with the structural scenario is from the unconditional distribution of the structural shocks (i.e., from the standard normal distribution). We will use the Kullback-Leibler (KL) divergence, D_{KL}, as a measure of how different the two distributions of structural shocks are.[11] Specifically, $D_{\mathrm{KL}}(P\|Q) = \int_X p \log(\frac{p}{q})\,\mathrm{d}\mu$. 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). Denote with $\mathcal{N}_{SS}$ the distribution of the structural shocks compatible with the structural scenario and $\mathcal{N}_U$ the unconditional distribution of structural shocks. In our case, since the unconditional distribution of the shock is a standard normal distribution, we have that the KL divergence between $\mathcal{N}_U$ and $\mathcal{N}_{SS}$ is

$$D_{\mathrm{KL}}(\mathcal{N}_U\|\mathcal{N}_{SS}) = \frac{1}{2}\left(\mathrm{tr}\left(\boldsymbol{\Sigma}_\varepsilon^{-1}\right) + \boldsymbol{\mu}_\varepsilon'\boldsymbol{\Sigma}_\varepsilon^{-1}\boldsymbol{\mu}_\varepsilon - nh + \ln\left(\det\boldsymbol{\Sigma}_\varepsilon\right)\right) \tag{18}$$

where $\boldsymbol{\mu}_\varepsilon$ and $\boldsymbol{\Sigma}_\varepsilon$ are the mean and variance of the shocks under the structural scenario given by Equations (13) and (14).

While it is straightforward to compute the KL divergence between $\mathcal{N}_{SS}$ and $\mathcal{N}_U$ using

[11] Since the KL divergence is invariant to linear transformations the KL divergence of the structural shocks from the standard normal distribution is equivalent to the divergence of the distribution of the conditional forecast from the distribution of the unconditional forecast.

Equation (18), it is difficult to grasp whether any value for the KL divergence is large or small. In other words, although the KL divergence can easily be used to evaluate whether structural scenario A is further away from the unconditional forecast than structural scenario B, it is hard to say *how* far away they are from from the unconditional forecast.

To ease the interpretation of the KL divergence, McCulloch (1989) proposes "calibrating" the KL divergence between two generic distributions P and Q using the KL divergence between two easily interpretable distributions. In particular, he suggests calibrating the KL divergence between two generic distributions P and Q to a parameter q that would solve the following equation

$$D_{\mathrm{KL}}(\mathrm{Bern}(0.5) \| \mathrm{Bern}(q)) = D_{\mathrm{KL}}(\mathcal{N}_U \| \mathcal{N}_{SS}),$$

where $\mathrm{Bern}(q)$ is a Bernoulli distribution with probability q. Hence, the calibrated parameter q maps the KL divergence between two generic distributions P and Q to the distance between two Bernoulli distributions, one with probability q and the other with probability 0.5. It can be shown that $q \in [0.5, 1]$. In this way, 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 structural shocks under the structural scenario considered is not at all far from the unconditional distribution of the shocks. Therefore, the structural scenario considered is quite realistic. Similarly, with a value of $q = 0.99$ the structural scenario requires a substantial deviation of the structural shocks from their unconditional distribution, suggests that the scenario is extreme and therefore quite unlikely.

A drawback of McCulloch's (1989) approach in our setting is that the probability q is not scale invariant. Specifically, it increases quickly to one as nh, the dimension of the structural scenario, increases.[12] To solve this problem, we propose using two binomial distributions instead of two Bernoulli distributions. Let $\mathcal{B}(m, p)$ denote the binomial distribution that runs

[12] In fact, it is easy to show from Equation (18) that, for any $\boldsymbol{\mu}_\varepsilon \neq \mathbf{0}$ and/or $\boldsymbol{\Sigma}_\varepsilon \neq \mathbf{I}$, the KL divergence between $\mathcal{N}_{SS}$ and $\mathcal{N}_U$ increases linearly with nh. Thus $q \to 1$ for any structural scenarios with either n, the number of variables, or h, the forecast horizon, big enough.

m independent experiments, each of them with probability p of success. The parameter m allows us to control for the dimension of the problem, effectively scaling the KL divergence by the dimension of the structural scenario. Hence, as before, we suggest calibrating the KL divergence between two generic distributions P and Q to a parameter q that would solve the following equation

$$D_{\mathrm{KL}}(\mathcal{B}(nh, 0.5)\|\mathcal{B}(nh, q)) = D_{\mathrm{KL}}(\mathcal{N}_U\|\mathcal{N}_{SS}).$$

The solution to the equation is

$$q = \frac{1 + \sqrt{1 - e^{-\frac{2z}{nh}}}}{2}, \tag{19}$$

where $z = D_{\mathrm{KL}}(\mathcal{N}_U\|\mathcal{N}_{SS})$. The interpretation of q remains in line with McCulloch's (1989) original idea.[13]

4 Algorithms for structural scenario analysis

In this section we develop algorithms to implement the structural scenario analysis. Specifically, we extend the Gibbs sampler algorithm in Waggoner and Zha (1999) to implement the structural scenario analysis in set and partially identified SVARs. This method can be easily extended to exactly and fully identified SVARs. The algorithm can also be used to implement conditional-on-variables and conditional-on-shocks forecasts as special cases. We use a uniform-normal-inverse-Wishart distribution over the orthogonal reduced-form parameterization as defined in Arias, Rubio-Ramirez, and Waggoner (2016b). Hence, the posterior distribution is also a uniform-normal-inverse-Wishart distribution over the orthogonal reduced-form parameterization, as it is extremely easy to generate draws from it.

To simplify both the notation and the exposition, when presenting the algorithms we only focus on traditional sign restrictions, as in Uhlig (2005), Canova and Nicolo (2002) and Rubio-Ramirez, Waggoner, and Zha (2010). The algorithms can also be easily modified to

[13] If we set $nh = 1$ in Equation (19), we obtain the calibrated q for McCulloch's (1989) original idea: $q = (1 + \sqrt{1 - e^{-2z}})/2$.

implement the recently proposed narrative sign restrictions as in Antolin-Diaz and Rubio-Ramirez (2016) and zero restrictions by using the methods described in Arias, Rubio-Ramirez, and Waggoner (2016b).

Let $\mathbf{S}_j$ be an $s_j \times r$ matrix of full row rank, where $0 \leq s_j$ where the $\mathbf{S}_j$ will define the traditional sign restrictions on the j^{th} structural shock for $1 \leq j \leq n$. In particular, we assume that if $(\mathbf{A}_0, \mathbf{A}_+)$ satisfy the restrictions, then

$$\mathbf{S}_j \mathbf{F}(\mathbf{A}_0, \mathbf{A}_+)\mathbf{e}_j > \mathbf{0} \text{ for } 1 \leq j \leq n,$$

where $\mathbf{e}_j$ is the j^{th} column of $\mathbf{I}_n$.

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

1. *Conditioning on* $\mathbf{y}^{T+h,(i-1)} = [\mathbf{y}^T, \mathbf{y}^{(i-1)}_{T+1,T+h}]$*, 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* $\mathbf{S}_j\mathbf{F}(f_h^{-1}(\mathbf{B}^{(i)}, \mathbf{\Sigma}^{(i)}, \mathbf{Q}^{(i)}))\mathbf{e}_j > \mathbf{0}$ *for* $1 \leq j \leq n$*, otherwise return to Step 1*

4. *Conditioning on* $(\mathbf{B}^{(i)}, \mathbf{\Sigma}^{(i)}, \mathbf{Q}^{(i)})$ *and* $\mathbf{y}^T$*, draw* $y^{(i)}_{T+1,T+h}$ *using Equation (9).*

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

The natural initialization can be done by using Equation (9) and the peak of the likelihood function or even a random draw from the posterior. Note that Algorithm 1 can be quite inefficient as it discards the draws, $(\mathbf{B}^{(i)}, \mathbf{\Sigma}^{(i)})$, if the associated orthogonal matrix, $\mathbf{Q}^{(i)}$, does not satisfy the restrictions. A more efficient version of the algorithm can be considered as follows.

Algorithm 2. *Initialize* $\mathbf{y}^{T+h,(0)} = [\mathbf{y}^T, \mathbf{y}^{(0)}_{T+1,T+h}]$.

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

2. *For each draw* $(\mathbf{B}^{(i,k)}, \mathbf{\Sigma}^{(i,k)})$, *make* M *draws* $\mathbf{Q}^{(i,k,m)}$ *independently from the uniform distribution over the set of orthogonal matrices.*

3. *Retain the triplets* $(\mathbf{B}^{(i,k)}, \mathbf{\Sigma}^{(i,k)}, \mathbf{Q}^{(i,k,m)})$ *from the set of triplets that satisfy the restrictions* $\mathbf{S}_j \mathbf{F}(f_h^{-1}(\mathbf{B}^{(i,k)}, \mathbf{\Sigma}^{(i,k)}, \mathbf{Q}^{(i,k,m)}))\mathbf{e}_j > \mathbf{0}$ *for* $1 \leq j \leq n$.

4. *Choose randomly a triplet* $(\mathbf{B}^{(i,k)}, \mathbf{\Sigma}^{(i,k)}, \mathbf{Q}^{(i,k,m)})$ *from the set obtained in Step 3, and call it* $(\mathbf{B}^{(i)}, \mathbf{\Sigma}^{(i)}, \mathbf{Q}^{(i)})$.

5. *Conditioning on* $(\mathbf{B}^{(i)}, \mathbf{\Sigma}^{(i)}, \mathbf{Q}^{(i)})$ *and* $\mathbf{y}^T$, *draw* $y^{(i)}_{T+1,T+h}$ *using Equation (9).*

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

It is also worth noticing that, owing to the independence of the K draws of the reduced-form parameters, step 1 of Algorithm 2 can be parallelized, so as to further increase the computational efficiency of the algorithm.

4.1 The importance of using all available identification restrictions

Note that when using restrictions that set identify the model, the results of the structural scenario analysis will be robust across the set of structural models that satisfy the restrictions. This attractive feature will come at the cost of very wide confidence bands around the forecast. Most important, there is the risk of including many structural models with implausible implications for elasticities, structural parameters, shocks and historical decompositions. This point has been forcefully argued by Kilian and Murphy (2012), Arias, Caldara, and Rubio-Ramirez (2016a), and Antolin-Diaz and Rubio-Ramirez (2016).

Unlike unconditional forecasting and conditional-on-observables forecasting, structural scenario analysis requires the structural parameters to be identified. For instance, in the applications with monetary policy shocks that we present below, we find that a strategy based exclusively on the traditional sign restrictions on impulse response functions often leads to implausibly large elasticities of observable variables to a monetary policy shock, mirroring the results of Kilian and Murphy (2012).[14] In the examples that follow, we will propose using a combination of traditional sign and zero restrictions at various horizons, as in Arias, Caldara, and Rubio-Ramirez (2016a), and the recently proposed narrative sign restrictions of Antolin-Diaz and Rubio-Ramirez (2016), to narrow down the set of admissible structural parameters and obtain meaningful structural scenarios.

5 Examples

We now propose two examples to illustrate the methods we have described above. In particular, we will first analyze some monetary policy structural scenarios. The second example will be about some structural scenarios related to possible stress-test analysis. Both models are set identified, but the first one is fully identified, while the second is partially identified.

5.1 Monetary policy structural scenarios

Let us begin analyzing some monetary policy structural scenarios. We will consider a model with three variables: the quarterly growth rate of real GDP, the quarterly growth rate of the core PCE deflator, and the federal funds rate at quarterly frequency, from 1955 to 2015. We consider five lags and the Minnesota prior over the reduced-form parameters. We will first compare the results of unconditional forecasting, conditional-on-observables forecasting, and structural scenario analysis. Identification of the structural parameters is only necessary for

[14] These elasticities are much larger than the upper bound reported by 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.

the latter. However, by identifying the structural parameters, we will be able to understand and interpret the results in light of their implications for the structural shocks even for unconditional and conditional-on-observables forecasting. In the last part of this section on monetary policy structural scenarios, we will use our structural scenario analysis methods to compare some policy alternatives that we think the FOMC had as of December 2015. The intention of such an exercise is to show how our structural scenario analysis can help in the policy debate. In this last part, we will also show how important it is to consider uncertainty around the conditioning path for observables.

We identify three structural shocks: a monetary policy (MP) shock, an aggregate demand (AD) shock, and an aggregate supply (AS) shock. We identify the structural shocks using zero and traditional sign restrictions on the IRFs and narrative sign restrictions. In particular, we use the zero and traditional sign restrictions on IRFs displayed in Table 1. First, a MP shock reduces output and inflation and increases the federal funds rate on impact and it is restricted to have zero long-run impact on the level of output. Second, a contractionary AD shock reduces output, inflation, and the interest rate on impact, and is restricted to have a zero long-run impact on the level of output. Third, an AS shock is restricted to reduce real GDP and increase inflation and the nominal interest rate on impact. Taken together, our identification scheme implies that the AS shock is the only one with a permanent effect on the level of output, as in Blanchard and Quah (1989), and Gali (1999). These restrictions are consistent with a wide class of standard New Keynesian models and are widely used in the SVAR literature, including studies by Bernanke and Mihov (1998), Erceg, Guerrieri, and Gust (2005), and Kilian and Lutkepohl (2017).

As mentioned above, traditional sign and zero restrictions are usually not sufficient to rule out many structural models with implausible implications for the structural parameters. We therefore follow Antolin-Diaz and Rubio-Ramirez (2016) and impose in addition the following narrative sign restrictions.

Table 1: Traditional Sign and Zero Responses

	Impact			Long Run		
Variable / Shock	MP	AD	AS	MP	AD	AS
Real GDP	−	−	−	0	0	
Core PCE inflation	−	−	+			
Federal funds rate	+	−	+			

Note: The long-run restriction is implemented at the infinite-horizon cumulative IRF of output growth, and at the 32-quarter horizon for the level of the inflation rate.

Narrative Sign Restriction 4.1.1. *The monetary policy shock for the observation corresponding to the fourth quarter 1979 must be of positive value.*

Narrative Sign Restriction 4.1.2. *For the observation corresponding to the fourth quarter of 1979, a monetary policy shock is the overwhelming driver of the unexpected movement in the federal funds rate. In other words, the absolute value of the contribution of monetary policy shocks to the unexpected movement in the federal funds rate is larger than the sum of the absolute value of the contributions of all other structural shocks.*

These two narrative sign restrictions have been shown to be very useful to help identify monetary policy shocks.

For the results that follow, it will be useful to examine the forecast error variance decomposition resulting from our identification scheme, shown in Figure 2. The figure makes it clear that at horizons greater than one year, aggregate demand shocks are the primary driver of unexpected movements in the federal funds rate. In other words, the bulk of the unexpected variation in interest rates is due to the systematic response of the monetary authority to aggregate demand shocks. As we will see below, these results have important implications for the typical conditional forecasts restricting the path of the interest rate over the forecast horizon.

Figure 2: MONETARY POLICY: FORECAST ERROR VARIANCE DECOMPOSITION

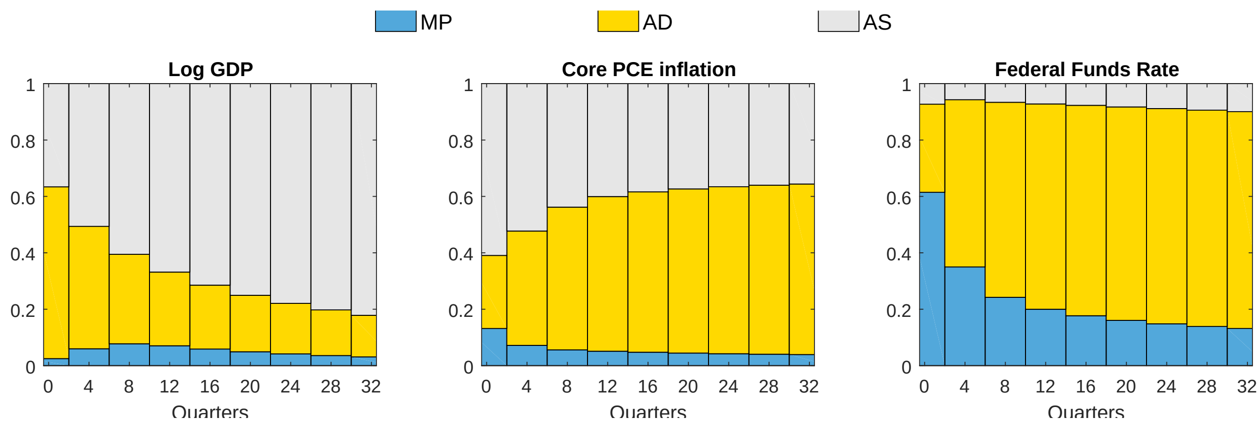

Note: The figure shows the mean posterior forecast error variance decomposition. For each panel, the colored bars represent the fraction of the total variance of the respective endogenous variable attributable to a specific structural shock at the horizon given by the horizontal axis.

5.1.1 Unconditional forecasting

Figure 3 considers the unconditional forecast of the model, as in Section 2.1. The unconditional forecast foresees that output growth will increase slightly and stay in the vicinity of 2 percent, that inflation will recover gradually toward 2 percent and that the federal funds rate increases in a very gradual manner, approaching 2.5 percent by the end of the forecast horizon.

Although identification of the structural parameters is not required to produce an unconditional forecast, we can interpret the forecast through the lens of the structural identification. Panel (b) displays the probability density function (PDF) of the structural shocks implied by this unconditional forecast. Each of the dashed PDFs represents the density of the estimated structural shocks at $t = T + 1 \ldots T + 12$. The PDF of the unconditional distribution of the shocks, i.e., the standard normal distribution, is represented as a gray shaded area. As expected, the future structural shocks will be normally distributed with mean zero and unit variance. This is because the unconditional forecast reflects no information about the future structural shocks beyond their unconditional distribution.

Figure 3: Monetary Policy: Unconditional Forecast

(a) Unconditional Forecasts

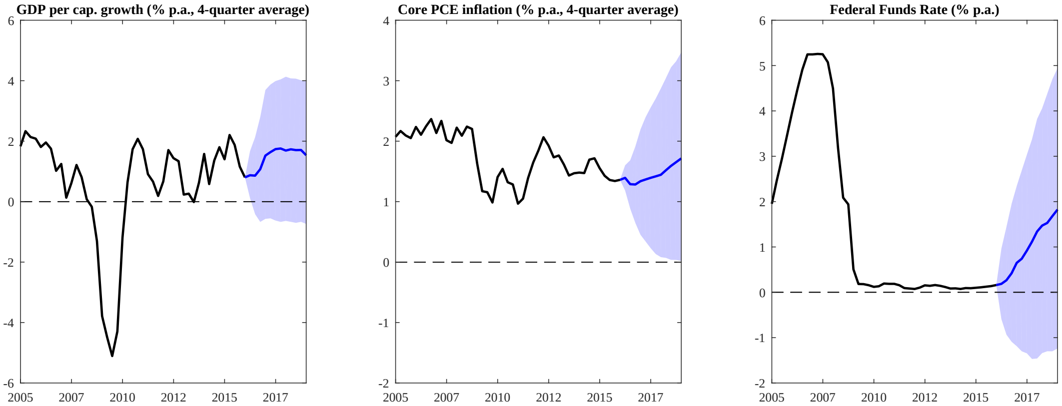

(b) Unconditional Structural Shocks

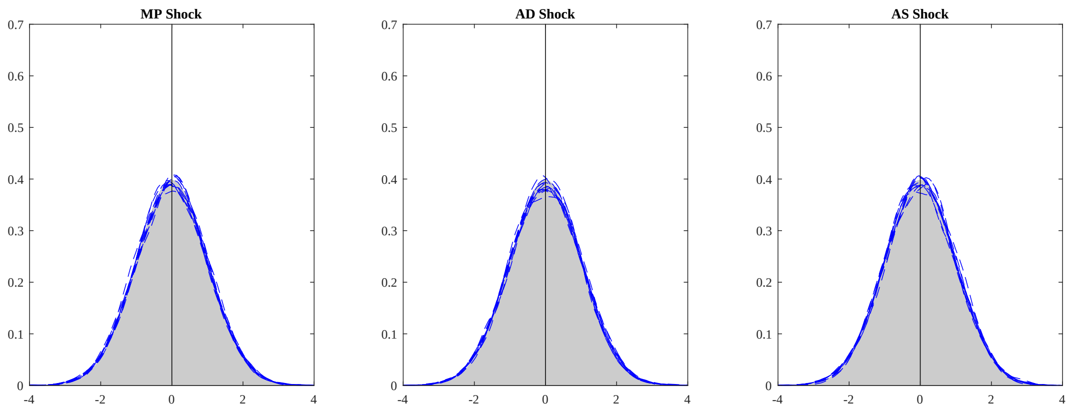

Note: In the top panel, for each column, the solid black lines represent actual data, the solid blue line is the median forecast and the blue shaded areas denote the 68 percent pointwise credible sets around the forecasts. The lower panel displays the PDFs of the structural shocks implied by the forecast for every $t = T + 1 \ldots T + 12$. The gray shaded area is the PDF of a standard normal distribution.

5.1.2 Conditional-on-observables forecasting

We now assume that we want to condition on the future path of the federal funds rate. We assume that from $t = T + 1 \ldots T + 12$ the federal funds rate increases by 50 basis points each quarter until it reaches 525 basis points. This pace of tightening of the interest rate is identical to the one observed in the mid-2000s but substantially faster than the model's unconditional forecast. Figure 4 considers the conditional-on-observables forecast.[15] As can be seen from the figure, the conditional-on-observables forecast foresees that inflation is increasing rapidly, and output is experiencing a boom, compared to the unconditional forecast.

Once again, although identification of the structural parameters is not required to produce a conditional-on-observables forecast, interpreting this type of conditional forecast through the lens of the structural identification can shed light on the economic intuition behind the results. Panel (b) displays the probability density function (PDF) of the structural shocks implied by this conditional forecast. As before, each of the dashed PDFs represents the density of the estimated structural shocks at $t = T + 1 \ldots T + 12$ and the PDF of the unconditional distribution of the shocks, i.e., the standard normal distribution, is represented as a gray shaded area. It is clear from the PDFs that the conditional forecast entails a combination of small positive (i.e., contractionary) monetary policy shocks and negative (i.e., expansionary) aggregate demand shocks.

These results, taken together with the forecast error variance decompositions of Figure 2, allow us to understand the conditional forecast of output and inflation. The given path for the federal funds rate implies a persistent unexpected increase in the interest rate that lasts for three years. At this horizon, aggregate demand shocks are the most important driver of the federal funds rate. Therefore, the conditional forecast reflects the fact that the most likely shock to have caused such an increase in the interest rate is an expansionary aggregate demand shock that is also increasing output and inflation.

[15] All the results that follow are the result of implementing 5000 draws of Algorithm 2 above, of which the first 1000 are discarded as burn-in draws.

Figure 4: Monetary Policy: Conditional-on-Observables Forecast

(a) Conditional-on-Observable Forecasts

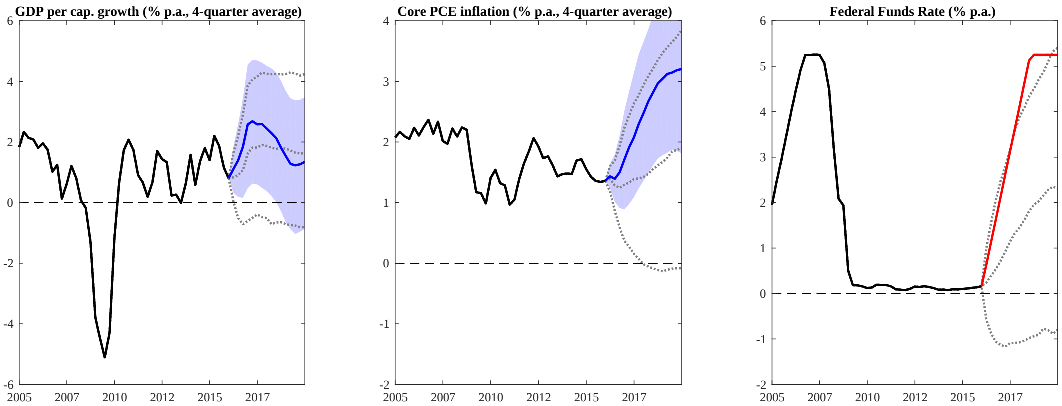

(b) Conditional-on-Observable Structural Shocks

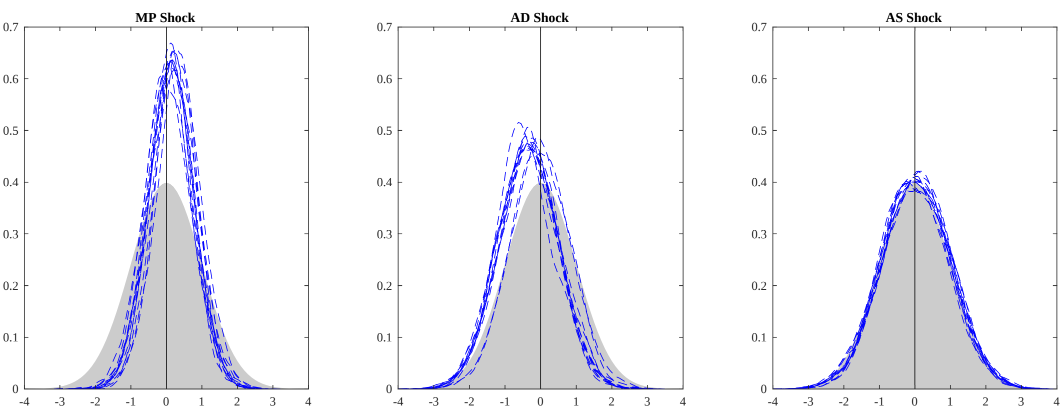

Note: In the top panel, for each column, the black solid lines represents actual data, the solid red line is the conditioning assumption on the observables, the solid blue line is the median forecast for the unrestricted variables and periods, and the blue shaded areas denote the 68 percent pointwise credible sets around the forecasts. The dotted black lines represent the median and contours of the 68 percent credible set around the unconditional forecast. The lower panel displays the PDFs of the structural shocks implied by the forecast for every $t = T+1 \ldots T+12$. The gray shaded area is the PDF of a standard normal distribution.

Figure 5: Monetary Policy: Structural Scenario

(a) Structural Scenario Forecasts

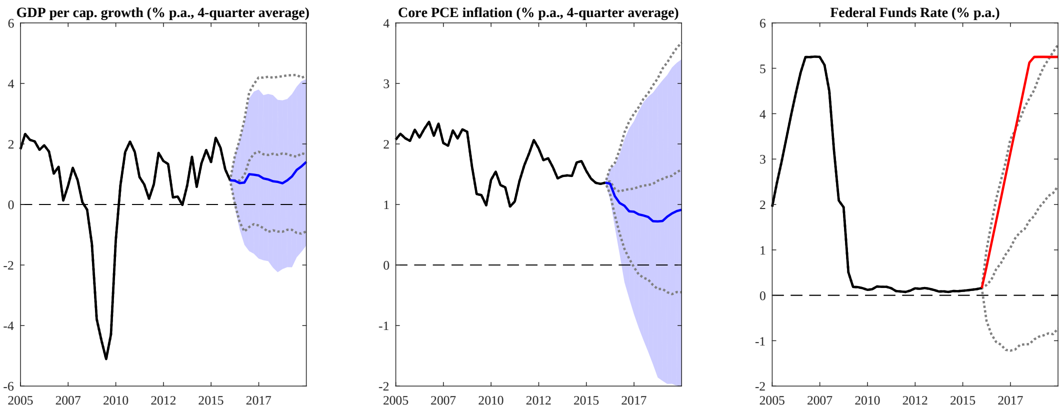

(b) Structural Scenario Shocks

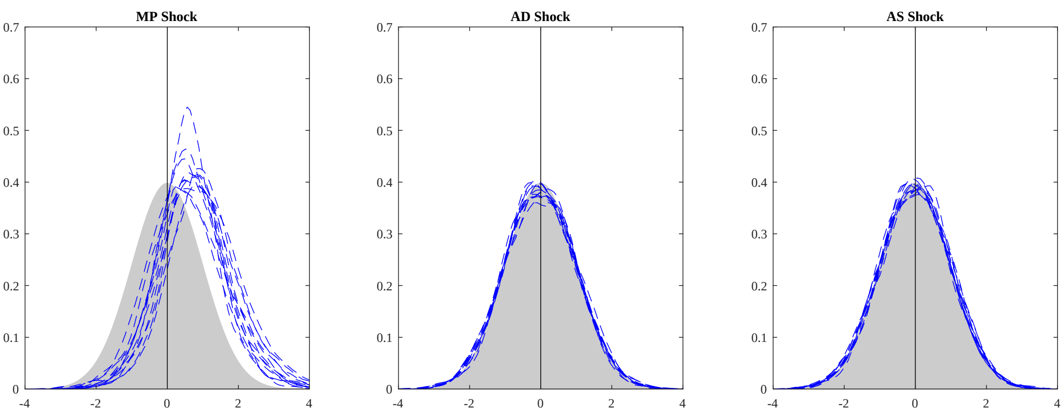

Note: In the top panel, for each column, the solid black lines represent actual data, the solid red line is the conditioning assumption on the observables, the solid blue line is the median forecast for the unrestricted variables and periods, and the gray shaded areas denote the 68 percent pointwise credible sets around the forecasts. The dotted black lines represent the median and contours of the 68 percent credible set around the unconditional forecast. The lower panel displays the PDFs of the structural shocks implied by the forecast for every $t = T + 1 \ldots T + 12$. The gray shaded area is the PDF of a standard normal distribution.

5.1.3 Structural scenario analysis

We now use the results of Section 3.4 to analyze the following structural scenario. As in the previous subsection, the federal funds rate increases by 50 basis points each quarter until it reaches 525 basis points. However, here we impose the restriction that the monetary policy shock is the key driver of this structural scenario. In other words, the AD and AS shocks are restricted to retain their unconditional distributions.

Panel (a) of Figure 5 shows that the results are strikingly different from conditional-on-observables forecasting: inflation falls below 1 percent at the end of the forecast horizon, and output slows down. As can be seen in Panel (b), by construction the monetary policy shock is the only one that deviates from its unconditional distribution. Since the monetary policy shock is a less important driver of the federal funds rate at business cycle frequencies, a larger sequence of contractionary monetary policy shocks is required to produce the given path of the funds rate. This larger sequence of contractionary monetary policy shocks exerts a strong negative impact on output and inflation. Consequently, the resulting forecast for output and inflation is weaker than the unconditional forecast.

These results highlight our main point: conditional-on-observables forecasting is equivalent to asking the model what combination of structural shocks is on average more likely to have generated the given path for the conditioning variable. In that case, the methods of Waggoner and Zha (1999) give the appropriate answer. But in many instances, the researcher might be interested not in tracing the effects of the average combination of structural shocks, but in conditioning on a particular structural shock driving the forecast. In which case, the methods described in Section 3.4 must be used. The two approaches will often give substantially different results. The two approaches can therefore be regarded as complementary, depending on the question to be answered.

It is worth noticing that the 68 percent high posterior density bands around the median forecasts are substantially wider in the case of the structural scenario analysis. As mentioned above, the conditional-on-observables forecast is invariant to the structural identification, and therefore, the only uncertainty surrounding the forecast is the uncertainty about the reduced-form parameters. On the contrary, when the structural scenario is considered, the uncertainty about the structural parameters must be taken into account.

5.1.4 Comparison of policy alternatives

We will now use our structural scenario analysis methods to compare some policy alternatives that we think the FOMC had as of December 2015. We choose December 2015 because that was the moment that the Federal Reserve started increasing the federal funds rate. Figure 6 analyzes three structural scenarios for the path of the interest rate. For the three paths we consider in Figure 6, we assume that there is no uncertainty around the interest rate path. It is also the case that, for all of them, we impose the restriction that the monetary policy shock is the key driver of the three structural scenarios, i.e., that the AD and AS shocks are restricted to retain their unconditional distributions.[16]

The first of the three structural scenarios in Figure 6 is the "Baseline SEP." This structural scenario is borrowed from the Summary of Economic Projections published by the FOMC.[17] It foresees the federal funds rate increasing by 25 basis points each quarter until it reaches 3.5 percent. As can be seen from Panel (a), this structural scenario is associated with subdued inflation, which hovers just above 1 percent for the forecast horizon. Next, we consider the first alternative structural scenario, which we denote "lower for longer," seen in Panel (b). In this structural scenario, the federal funds rate is kept at zero for an additional two years, after which it is increased by 50 basis points every quarter over the forecast horizon. As can be seen from the figure, this structural scenario is associated with a modest boom in output and a faster return of inflation to the 2 percent target. Finally, we consider the "Tighter" structural scenario of the previous subsection, in which the fed funds rate is increased 50

[16] That is, in all the structural scenarios, the AD and AS shocks are constrained to their unconditional distribution.

[17] See Federal Open Market Committee, "*Summary of Economic Projections*," December 16, 2015.

Table 2: Plausibility of the Structural Scenarios

	No uncertainty		With uncertainty	
	KL divergence	Calibrated q	KL divergence	Calibrated q
Baseline SEP	5.57×10^4	1	4.45	0.70
Lower for Longer	5.92×10^4	1	4.69	0.71
Tighter	1.01×10^5	1	5.08	0.72

Note: We report the mean of the KL divergence across draws of the posterior. The distribution of the KL is narrowly centered around the mean for all cases, and is available upon request.

basis points every quarter until reaching 5.25 percent. As can be seen from Panel (c), this structural scenario is associated with low output and low inflation.

Table 2 looks at the plausibility of these alternative structural scenarios using the Kullback-Leibler divergence and its calibration as proposed in Section 3.5. The first thing to notice is that when the structural scenario is imposed with no uncertainty around the path for the fed funds rate, i.e., $\mathbf{\Omega}_f = \mathbf{0}$, as in Figure 6, the KL divergence is very high, on the order of tens of thousands. The calibrated q is equal to 1 for all three cases, deeming the three structural scenarios as highly unlikely. The reason for this is that the structural shocks need to be distorted a great deal from the unconditional distribution in order to produce a fixed path for one of the observables. The KL divergence and its calibrated q confirm the intuition in Andersson, Palmqvist, and Waggoner (2010), who suggest that conditioning on a fixed path for a particular variable and ignoring the uncertainty about the conditioning assumption can lead to unrealistic density forecasts. In Figure 7 we instead set the variance around the fed funds rate to its unconditional variance, $\mathbf{\Omega}_f = \overline{\mathbf{DD}}'$, as proposed by Andersson, Palmqvist, and Waggoner (2010). In this case, we obtain a more plausible structural scenario, as seen in the last two columns of Table 2. As expected, the uncertainty bands around the forecasts of the other variables are now wider. The values of the calibrated qs suggest that, whereas the baseline SEP forecast represents the most likely structural scenarios among the ones

Figure 6: Comparison of Policy Alternatives (No uncertainty)

(a) Baseline SEP

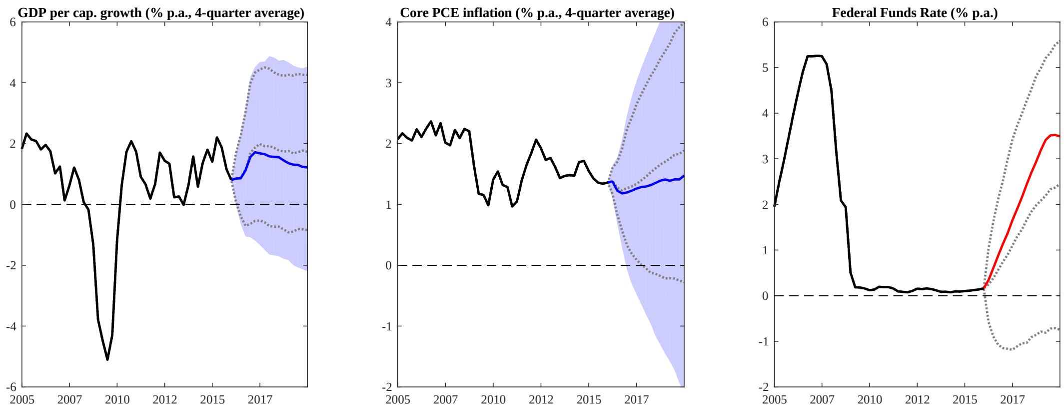

(b) Lower for Longer

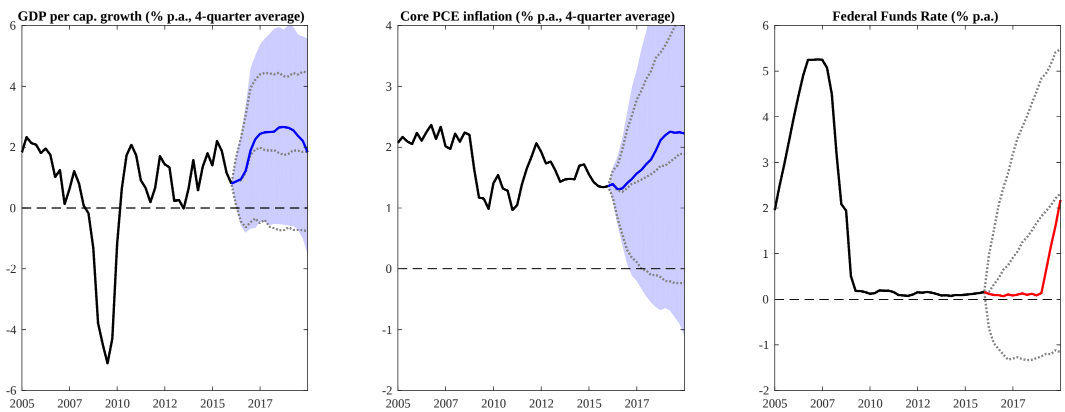

(c) Tighter

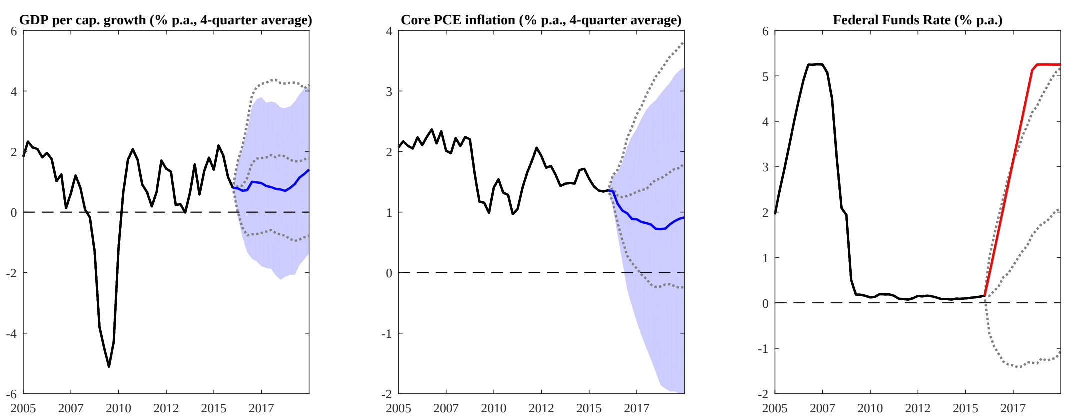

Figure 7: Comparison of Policy Alternatives (With uncertainty)

(a) Baseline SEP

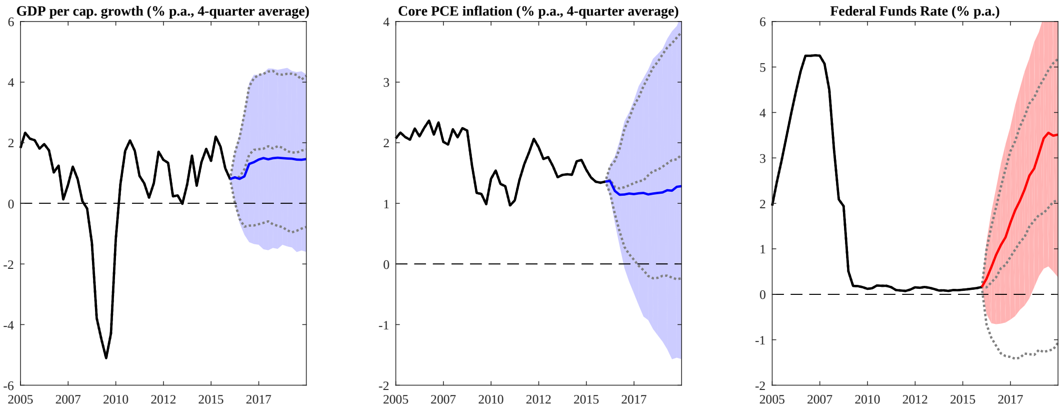

(b) Lower for Longer

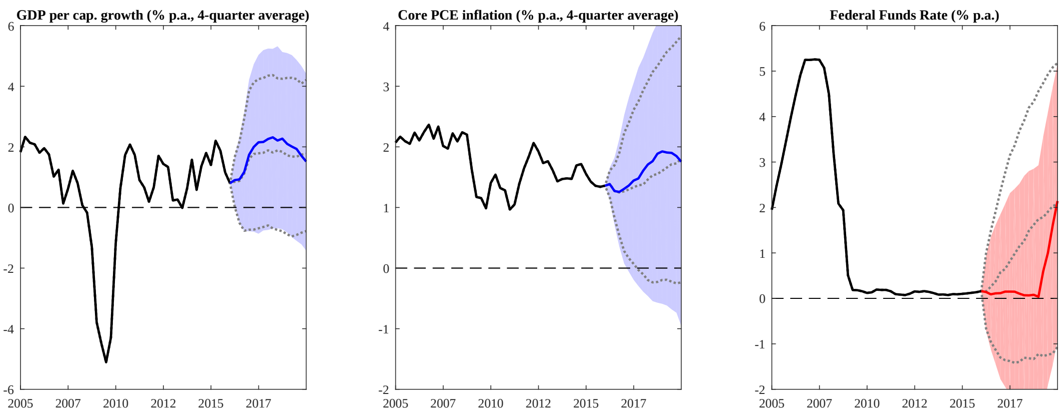

(c) Tighter

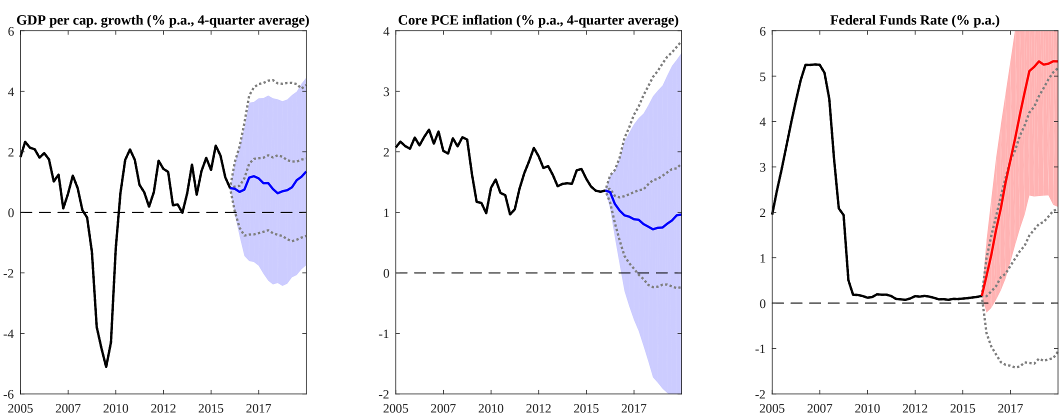

considered, the difference between the three is minimal and all of them are associated with important distortions in the unconditional distribution of the monetary policy shocks over the forecast horizon.

5.2 Stress-testing

Our second example considers a stress-testing exercise using a structural scenario analysis. In particular, we study the impact of an economic recession on asset prices and bank profitability. Our key objective is to highlight that the potential impact of the recessionary episodes can be very different depending on the main structural shock driving the recession. Specifically, we will show how a recession caused by financial structural 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. In order to make this point, we use a medium scale VAR and we identify a single structural shock: a financial shock. The VAR contains four key macroeconomic variables: the quarterly change in real GDP, the quarterly change in the core PCE deflator, the 3-month Treasury bill rate, and the unemployment rate; 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 in the banking system as a whole, the return on equity (ROE) of FDIC-insured institutions.[18] The data are quarterly and seasonally adjusted, from 1984 to 2016.

As in the previous subsection, the identification scheme employs a combination of traditional sign and narrative sign restrictions. On the traditional sign side, the financial

[18] The Haver mnemonics for the data are as follows: 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 (FBAA@USECON-FLTG@USECON), TED spread (C111FRED@OECDMEI - FTBS3@USECON), ROE (USARQ@FDIC). We use Haver's seasonal adjustment function whenever the original data are not seasonally adjusted.

Table 3: STRESS TEST: CONDITIONING PATH FOR GDP AND UNEMPLOYMENT

Variable / Period	1	2	3	4	5	6	7	8	9	10	11	12
Real GDP growth	2.0	1.8	1.5	0.7	-1.5	-2.8	-2	-1.5	-0.5	1	1.4	2.6
Unemployment rate	4.5	4.5	4.7	5.0	5.2	5.8	6.3	6.8	7.1	7.3	7.4	7.4

shock is restricted to have a negative impact on stock prices and bank profitability and to increase the BAA and TED spreads. In addition, we will narrow down the set of admissible models using narrative sign restrictions. Our main source will be the account in Bernanke (2015), which provides a detailed eyewitness description of the events surrounding the global financial crisis of the fall of 2008. In particular, Bernanke 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 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). Following Bernanke's assertion that a large contractionary financial shock in the fourth quarter of 2008 triggered a large increase in spreads, we impose the following narrative sign restrictions

Narrative Sign Restriction 4.2.1. *The financial shock for the observation corresponding to the fourth quarter of 2008 must be of positive value.*

Narrative Sign Restriction 4.2.2. *In the fourth quarter of 2008, the financial shock is the overwhelming driver of the unexpected movement in the TED spread and credit spread. In other words, the absolute value of the contribution of financial shocks to the unexpected movement in these variables is larger than the sum of the absolute value of the contributions of all other structural shocks.*

To construct the scenario, we borrow from the Federal Reserve's "2017 Supervisory

Scenarios for Annual Stress Tests Required under the Dodd-Frank Act Stress Testing Rules and the Capital Plan Rule."[19] In this document, the Federal Reserve lays out various scenarios that banks participating in the Fed's stress tests can use to assess the impact of the scenarios on their loan books. We take the path of GDP and the unemployment rate described in the Fed's "adverse scenario." This scenario describes a mild recession in which output falls for five consecutive quarters and then recovers gradually, whereas the unemployment rate increases until it reaches 7.4 percent. The exact path of these two variables is given in Table 3.[20] Conditional on the paths for GDP and unemployment, we consider two distinct structural scenarios. The first is the one in which the recession is *not* driven by the financial shock. We call this a non-financial recession. To implement this, we restrict the financial shocks to retain the $\mathcal{N}(0,1)$ distribution, and allow the other structural shocks to depart from their unconditional distribution. In other words, this structural scenario captures an economic downturn driven by shocks other than the financial shock. Figure 8 displays the results of the structural scenario. The dotted lines represent the unconditional forecast, and the solid blue lines and gray shaded areas are, respectively, the median and 68 percent high posterior density intervals around the structural scenario. Inflation, interest rates, and stock and house prices all drop, whereas the oil price and the credit spread are slightly higher. The TED spread increases at first, before declining. There is a mild decline in bank profitability as measured by the ROE.

Figure 9 instead considers the structural scenario in which the financial shock is driving the recession, and the rest of the shocks retain the $\mathcal{N}(0,1)$ distribution. We call this a financial recession. In this case, for identical paths of GDP and unemployment, the decline in inflation is more pronounced, whereas oil prices now decline and bond yields rise. There is now a large increase in spreads, and, importantly, bank profitability declines much more

[19]See https://www.federalreserve.gov/newsevents/pressreleases/files/bcreg20170203a5.pdf downloaded on May, 25, 2017.

[20]The stress test envisioned by the Fed has a very abrupt decline in GDP from about 2 percent to -1.5 percent within one quarter, which we phase in gradually during the first four quarters.

Figure 8: Stress Test with Structural Scenario: Non-Financial Recession

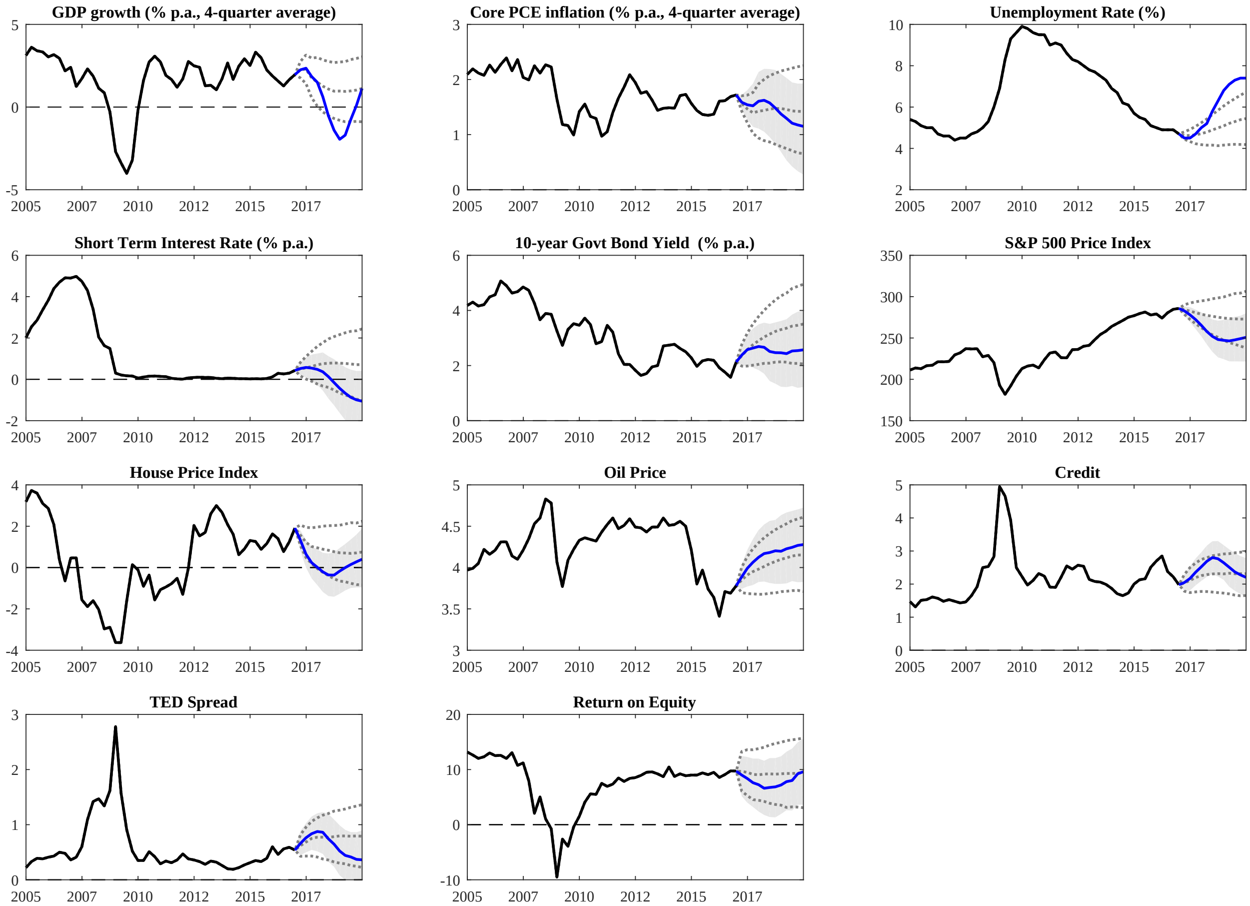

Note: For each panel, the dotted lines represent the unconditional forecast, and the solid blue lines and gray shaded areas are, respectively, the median and 68 percent high posterior density intervals around the scenario.

dramatically.

The exercise once again highlights the main point of considering structural scenarios: which structural shock is driving the restricted paths of observables can have important consequences for the implied paths of the unrestricted observables. In the case of stress testing, it is clear that not all recessions are alike, and which structural shock is driving the recession will matter.

Figure 9: Stress Test with Structural Scenario: Financial Recession

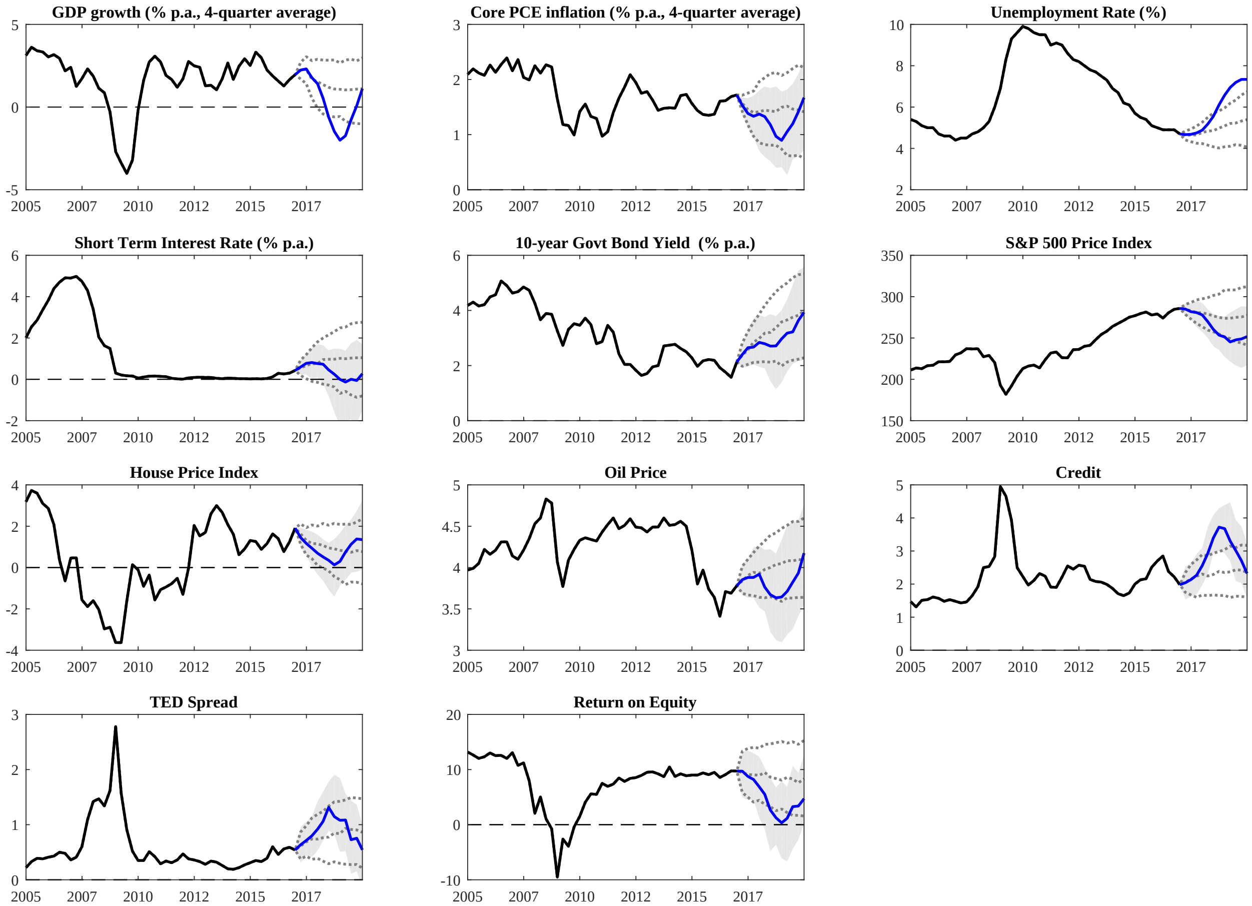

Note: For each panel, the dotted lines represent the unconditional forecast, and the solid blue lines and gray shaded areas are, respectively, the median and 68 percent high posterior density intervals around the scenario.

6 Conclusion

The assessment of conditional forecasts, i.e., quantifying likely future values of some macroeconomic variables given a hypothetical path for other variables, is an important part of the toolkit of applied macroeconomics. Conditional forecasting is usually a statistical exercise, involving the dynamic correlations among variables, and remaining silent about the underlying economic causes behind the forecast. However, more often than not, the researcher would like to analyze the conditional forecast through the lens of a structural model. This would allow her to assess the future value of some variables given a path for other variables that is driven by a specific set of structural shocks. We have called this exercise structural scenario analysis. Given an SVAR, the structural scenario analysis specifies a path for the future of some of the variables and which structural shocks are the drivers of the scenario. We provide tools to perform structural scenario analysis using SVARs. We develop Bayesian methods that can be used for set and partially identified models. This allows us to capture uncertainty about both the parameters and the structural model itself. In our illustrations, we show that for the same path of the conditioning variables, the result can be very different depending on which structural shock is assumed to drive the structural scenario. Hence, structural scenario analysis can be a useful complement to traditional conditional forecasting when an economic interpretation of the forecasts is sought.

A Vector notation for VAR forecasts

Assume that we want to forecast the observables for some period ahead using the VAR in Equation (3). 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

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

where $\mathbf{0}_{n\times n}$ is a $n \times n$ matrix of zeros. 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 $\boldsymbol{\Sigma} = (\mathbf{A}_0\mathbf{A}'_0)^{-1}$, $\mathbf{M}'\mathbf{M}$ depends only on the reduced-form parameters, $\mathbf{B}$ and $\boldsymbol{\Sigma}$, even though $\mathbf{M}$ depends on the structural parameters, $\mathbf{A}_0$ and $\mathbf{A}_+$.

B Details on the conditional-on-shocks restrictions to the structural shocks

In this Appendix we show that the shocks that are left unrestricted in a conditional forecast with conditional-on-shocks restrictions are drawn from the standard normal distribution in the solution for the conditional forecast. The distribution of $\varepsilon_{T+1,T+h}$ compatible with the conditional-on-shocks distribution of $\mathbf{y}_{T+1,T+h}$ is

$$\varepsilon_{T+1,T+h} \sim \mathcal{N}\left(\underline{\boldsymbol{\mu}}_{\varepsilon}, \underline{\boldsymbol{\Sigma}}_{\varepsilon}\right), \tag{B.1}$$

where

$$\underline{\boldsymbol{\mu}}_{\varepsilon} = \underline{\mathbf{D}}^* \left(\underline{\mathbf{C}}\mathbf{b}_{T+1,T+h} + \mathbf{g}_{T+1,T+h}\right) - \underline{\mathbf{D}}^*\underline{\mathbf{C}}\mathbf{b}_{T+1,T+h} = \underline{\mathbf{D}}^*\mathbf{g}_{T+1,T+h} \tag{B.2}$$

$$\underline{\boldsymbol{\Sigma}}_{\varepsilon} = \underline{\mathbf{D}}^*\boldsymbol{\Omega}_g \left(\underline{\mathbf{D}}^*\right)' + \hat{\underline{\mathbf{D}}}'\hat{\underline{\mathbf{D}}} \tag{B.3}$$

and $\underline{\mathbf{D}}^*$ is the generalized inverse of $\underline{\mathbf{D}} = \underline{\mathbf{C}}\mathbf{M}' = \boldsymbol{\Xi}$ and $\hat{\underline{\mathbf{D}}}$ is any $(nh - k_s) \times nh$ such that its rows form an orthonormal basis for the null space of $\underline{\mathbf{D}}$. Without loss of generality, one can always reorder the structural shocks so that $\boldsymbol{\Xi}$ has a block structure, $\boldsymbol{\Xi} = \left[\mathbf{I}_{k_s}, \mathbf{0}_{k_s\times(nh-k_s)}\right]$ and the first k_s structural shocks are the restricted ones while the rest are kept unrestricted. It can be shown that $\underline{\mathbf{D}}^* = \left[\mathbf{I}_{k_s}, \mathbf{0}_{k\times(nh-k_s)}\right]'$ and $\hat{\underline{\mathbf{D}}} = \left[\mathbf{0}_{k_s\times(nh-k_s)}, \mathbf{I}_{k_s}\right]'$. Using Equation (B.2), we have that the mean of the restricted structural shocks is

$$\underline{\boldsymbol{\mu}}_{\varepsilon} = \begin{bmatrix} \mathbf{g}_{t+1,t+h} \\ \mathbf{0}_{nh-k_s} \end{bmatrix}. \tag{B.4}$$

Moreover, since

$$\underline{\hat{\mathbf{D}}}'\underline{\hat{\mathbf{D}}} = \begin{bmatrix} \mathbf{0}_{k_s \times k_s} & \mathbf{0}_{k_s \times (nh-k_s)} \\ \mathbf{0}_{(nh-k_s)\times k_s} & \mathbf{I}_{nh-k_s} \end{bmatrix} \text{ and } \underline{\mathbf{D}}^{*}\boldsymbol{\Omega}_g(\underline{\mathbf{D}}^{*})' = \begin{bmatrix} \boldsymbol{\Omega}_g & \mathbf{0}_{k_s \times (nh-k_s)} \\ \mathbf{0}_{(nh-k_s)\times k_s} & \mathbf{0}_{(nh-k_s)\times(nh-k_s)} \end{bmatrix},$$

the variance of the restricted structural shocks in (B.3) simplifies to

$$\underline{\boldsymbol{\Sigma}}_{\varepsilon} = \begin{bmatrix} \boldsymbol{\Omega}_g & \mathbf{0}_{k_s \times (nh-k_s)} \\ \mathbf{0}_{(nh-k_s)\times k_s} & \mathbf{I}_{nh-k_s} \end{bmatrix}. \tag{B.5}$$

From Equations (B.4) and (B.5) it is easy to see that the first k entries of $\varepsilon_{T+1,T+h}$ have the distribution implied by Equation (15), while the rest of the structural shocks have a standard normal distribution, i.e., their distribution is unaltered from the unconditional distribution.

C Alternative implementation using Kalman filtering

The existing literature suggests that the same results from the methods we describe in this paper can be obtained by using state-space methods. For example, Clarida and Coyle (1984) show that one can implement the Doan, Litterman, and Sims (1986) conditional point forecasts using the Kalman filter. Building on this insight, Banbura, Giannone, and Lenza (2015) suggest that state-space methods can also be used to obtain conditional-on-observables forecasts as in Section 3.2. We now show that this equivalent algorithm using the Kalman filter is applicable to the structural scenario analysis of Section 3.4 as well. The Kalman filter implementation has the additional advantage that it can easily handle any pattern of missing data and mixed frequencies in the variables.

C.1 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:

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

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

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:

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

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) and the state-space representation in Equations

(C.1)-(C.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., Carter and Kohn (1994)) 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, Giannone, and Lenza (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.[21]

The same procedure can be adapted for the case of the conditional structural scenario. Suppose that one wants to produce forecasts for s time periods into the future. Denote with $*$ an observation for which all data entries are missing, denote $(\hat{y}_{i,t}, *)'$ for $T+1 \leq t \leq T+s$ 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 $(\hat{\varepsilon}_{j,t}, *)'$ for $T+1 \leq t \leq T+s$ 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 (C.1)-(C.2), in which the measurement vector is constructed as $\tilde{\mathbf{y}}_t'$ in the following way:

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

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

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 s 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 $T+s$ is used to fill the corresponding missing observations, and the missing observations for the structural shocks

[21]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 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, Davis, and Salehi (1990)).

other than the shock of interest are filled in with draws from their unconditional distribution.[22] One can then draw $y^{(i)}_{T+1,T+h}$ using the Kalman smoother, as opposed to Equation (9), in Step 5 of Algorithm 2. These two procedures produce identical results.

C.2 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 (9)-(11) involve operations of high dimensional matrix objects and, as such, can be computationally expensive, whereas the procedure based on the Kalman filter works recursively and can potentially reduce the computational burden significantly for longer forecast horizon, or a large number of variables. Table C.1 investigates which of the two procedures is more efficient numerically for different hypothetical problem settings.[23] The baseline implementation of Algorithm 2 is faster in the majority of the conditional forecast computations. However, as T increases, drawing the conditional forecast using the Kalman smoother is more efficient, especially in a VAR with small lags.

The relationship between n, T, p and computing time, denoted t_e, is very different in each method, however. In the case of the baseline algorithm, as T rises, the increase in t_e from introducing higher lags or more variables is exponential. For example, t_e for $n = 3$ and $p = 1$ rises by a factor of 25 as T is increased. However, with $p = 1$ and $n = 11$, t_e rises by a factor of 362 as the forecast horizon rises from 1 to 60. The same is qualitatively true when keeping n fixed and increasing the number of lags, p. The implementation based on the use of the Kalman filter, on the other hand, does not exhibit any exponential increase in t_e as the dimensionality of the system increases. As the forecast horizon increases, t_e increases in a more linear fashion and is similar in magnitude across the p and n dimensions. For example, under a model of 3 variables and 1 lag, t_e increases by a factor of 1.23 as the forecast horizon (T) rises from 1 to 60. These times are also similar in

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

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

Table C.1: Computational efficiency of conditional forecast methods

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

magnitude as the number of lags increases. When the dimensionality rises due to an increase in n, t_e increases by a larger factor, but not in an explosive way as demonstrated by Algorithm 2.

While we find that the Kalman filter implementation is not more efficient than the baseline specification of Algorithm 2, the intuition of Banbura, Giannone, and Lenza (2015) remains intact –the recursive nature of the Kalman filter means that computational efficiency is not much affected as the forecast horizon increases, while for Equation (9), the high dimensionality of the data and long forecast horizon increase the computational burden significantly.[24] Therefore, there is a somewhat large fixed cost in initially computing a conditional forecast using the Kalman filter procedure, but once this has been done, the recursive nature of the procedure ensures that the additional computing requirements are minimal as the forecast horizon increases. On the other hand, with Equation (9)

[24]One reason why our result contradicts the ones in Banbura, Giannone, and Lenza (2015) is that Andersson, Palmqvist, and Waggoner's (2010) implementation of the conditional forecast is more efficient than the original procedure of Waggoner and Zha (1999).

there is a small fixed cost, but the computational requirements increase significantly as the forecast horizon increases.

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 (2016): "Narrative Sign Restrictions for SVARs," CEPR Discussion Papers 11517.

ARIAS, J. E., D. CALDARA, AND J. F. RUBIO-RAMIREZ (2016a): "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 (2016b): "Inference Based on SVARs Identified with Sign and Zero Restrictions: Theory and Applications," Forthcoming, *Econometrica.*

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

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

BERNANKE, B. S. AND I. MIHOV *(1998): "The liquidity effect and long-run neutrality,"* Carnegie-Rochester Conference Series on Public Policy*, 49, 149–194.*

BLANCHARD, O. J. AND D. QUAH *(1989): "The Dynamic Effects of Aggregate Demand and Supply Disturbances,"* American Economic Review*, 79, 655–673.*

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.*

Carter, C. K. and R. Kohn *(1994): "On Gibbs Sampling for State Space Models,"* Biometrika*, 81, 541–553.*

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.*

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

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.*

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

Erceg, C. J., L. Guerrieri, and C. Gust *(2005): "Can Long-Run Restrictions Identify Technology Shocks?"* Journal of the European Economic Association*, 3, 1237–1278.*

Gali, J. *(1999): "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?"* American Economic Review*, 89, 249–271.*

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

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.*

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.*

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

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

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.*

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.*