Documento de Trabajo - 2017/13

Structural Scenario Analysis and Stress Testing with Vector Autoregressions

Juan Antolín-Díaz
(Fulcrum Asset Management)

Ivan Petrella
(Warwick Business School)

Juan F. Rubio-Ramírez
(Emory University, Federal Reserve Bank of Atlanta & FEDEA)

fedea

Structural Scenario Analysis and Stress Testing with Vector Autoregressions*

Juan Antolín-Díaz
Fulcrum Asset Management

Ivan Petrella
Warwick Business School

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

Abstract

In the context of linear vector autoregressions (VAR), conditional forecasts and "stress tests" are typically constructed by specifying the future path of one or more endogenous variables, while remaining silent about the underlying structural economic shocks that might have caused that path. However, in many cases researchers are interested in choosing which structural shock is driving the path of the conditioning variables, allowing the construction of a "structural scenario" which can be given an economic interpretation. We develop efficient algorithms to compute structural scenarios, and show how this procedure can lead to very different, and complementary, results to those of the traditional conditional forecasting exercises. Our methods allow to compute the full posterior distribution around these scenarios in the context of set and incompletely identified structural VARs, taking into account parameter and model uncertainty. Finally, we propose a metric 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 bank profitability to an economic recession.

Keywords: Conditional forecasts, Bayesian methods, probability distribution, structural VARs
JEL Classification Numbers: C32, C53, E47.

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 macroeconomic variables change?"

*We are grateful to Gavyn Davies and Jonas Hallgren for helpful comments and suggestions.
†Corresponding author: Juan F. Rubio-Ramírez <juan.rubio-ramirez@emory.edu>, Economics Department, Emory University, Rich Memorial Building, Room 306, Atlanta, Georgia 30322-2240.

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 autoregressive (VAR) models, and Andersson et al. (2010) extend their results to the case when there is uncertainty about the paths of the conditioning variables.

These methods provide an answer which is based on the *reduced-form* representation of the VAR. In other words, they depend on the correlation structure of the data, and are agnostic about the underlying *structural* shocks that might be causing the movements in the variables. As an illustrative example, consider a three-variable VAR with output growth, inflation, and the policy interest rate. The details of the model are provided below in Section 5.1. A policy maker might want to use the 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?" The answer, computed using the methods of Waggoner and Zha (1999), is presented in Panel (a) in Figure 1. The dotted lines represent the median and 68% 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 three years is that inflation is about half a percent lower than in the unconditional forecast.

This might appear puzzling at first, as easy monetary policy is usually thought to stimulate output and inflation. However, one should keep in mind that if a large part of the movements in the federal funds rate represent the systematic reaction of the Fed to output and inflation developments, see, e.g. Leeper et al. (1996), the unconditional correlation between the interest rate and inflation observed in the data will be strongly positive. The conditional forecast 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?". If the federal funds rate is low, it is likely that the Fed is responding to some shock that has lowered output and inflation. In other words, the conditional forecast conveys information about what is the most likely combination of future structural shocks

Figure 1: Two Types of Forward Guidance in an SVAR

(a) Conditional forecasting

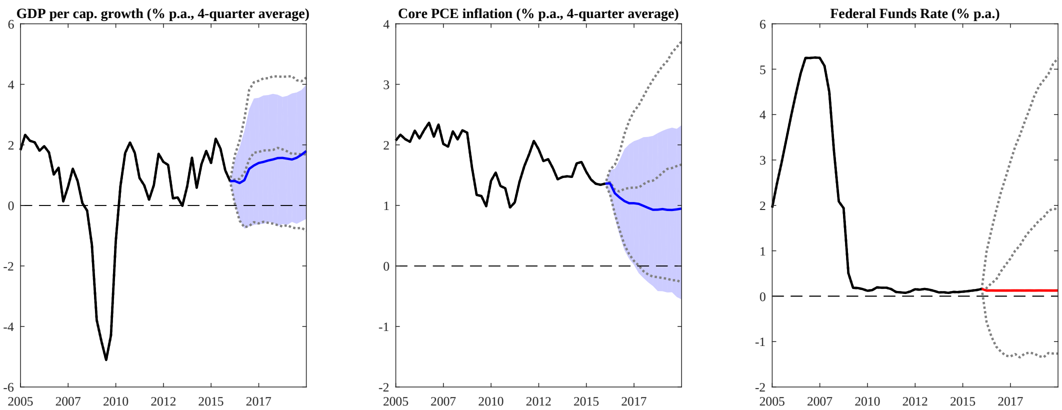

(b) Structural Scenario

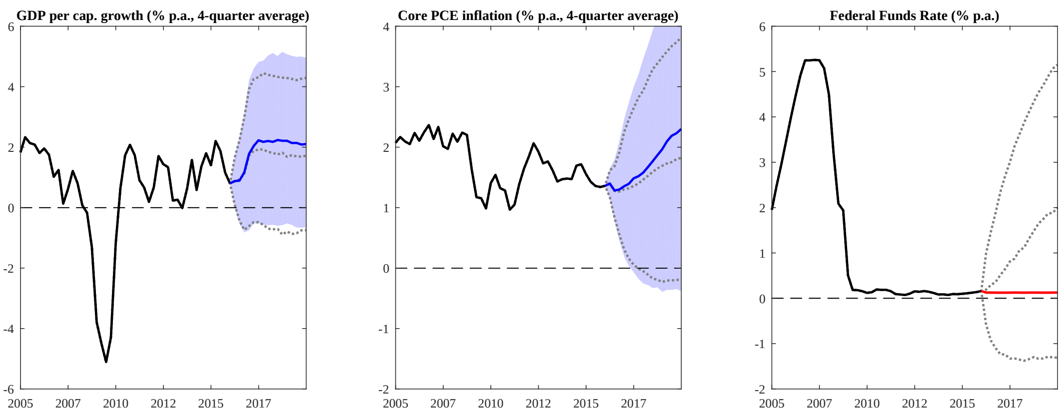

Note: For each column, the black solid lines represent actual data, the solid red line is the conditioning assumption on the observables, the solid blue line is the median forecast for the remaining variables and periods, and the blue shaded areas denote the 68 percent pointwise credible sets around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

that will generate the given path of the conditioning variable.[1]

In this paper, we develop tools to answer a question that differs from the one above in a subtle but important way. In many circumstances one might want to compute the conditional forecast to which one can give a structural interpretation, so that the path of the conditioning variable *reflects the effects of a particular structural shock*. Going back to the monetary policy example, the answer to the alternative question "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?" 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] We call this exercise "structural scenario analysis." The key difference with conditional forecasting is that the latter does not require to identify the underlying economic shocks in a Structural VAR (SVAR). On the contrary, structural scenario analysis requires going beyond the statistical correlations in the data and carefully thinking about the economic scenarios behind the exercise. As we will see, the results will critically depend on the identifying restrictions used to go from the reduced-form to the SVAR.

We are not the first to propose conditioning forecasts on structural shocks rather than on endogenous variables. Baumeister and Kilian (2014) perturb an identified SVAR of the oil market with exogenous shocks to trace out the impact on oil prices. This practice is also used routinely on the DSGE literature, where a fully specified structural model is also available, see Del Negro and Schorfheide (2013). Relative to these studies, our proposal carries several innovations. First, most previous studies specify conditions on structural shocks, which are unobserved. This is unintuitive and often requires iterating between the value of the shock and the desired movement in the endogenous variable (see, e.g., Clark and McCracken, 2014). Instead, we specify the scenarios in terms of paths for observables, allowing the model to calculate the likely path of the desired structural shock (or shocks) that generates such paths. Second we extend the existing Bayesian methods for computing the exact finite-sample distribution of conditional forecasts to the presence

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

[2] Similar to Campbell et al. (2012)'s "Odyssean" forward guidance, the exercise does not reveal any information about the 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.

of set and incompletely identified SVARs, whereas previous studies usually relied on a specific structural model. The methods developed in this paper fully take into account two sources of uncertainty that are typical of a structural VAR environment: the uncertainty about the structural parameters and uncertainty regarding the exogenous structural shocks over the forecast horizon.[3] Thus, our procedure is useful for researchers who do not want to fully specify a structural model, or do not completely trust the results of their model, and are interested in conclusions that are robust across a set of structural models. Lastly, we propose a way of assessing how plausible (or implausible) a conditional forecast is, effectively evaluating how far the conditional forecast is from the unconditional one. This tool offers a a simple way of comparing (and ranking) alternative scenario. Moreover it offers a simple metric to evaluate whether a particular scenario can be evaluated within a linear VAR setting (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, where 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 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 forecasting, whether on endogenous variables or structural shocks. Section 4 provides algorithms to implement our techniques, and to measure the plausibility of the different scenarios under consideration. Section 5 applies our techniques in two illustrations: a small, fully-identified

[3] The existing procedures for constraining the forecast to a particular path of structural shocks of the model (such as Baumeister and Kilian, 2014, and Clark and McCracken, 2014) have tended to disregard parameter uncertainty, whereas Waggoner and Zha (1999) find that ignoring parameter uncertainty can potentially results in misleading conditional forecasts.

New Keynesian VAR for analyzing the effects of a monetary policy tightening, and a larger, partially identified VAR 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 variables, $\boldsymbol{\varepsilon}_t$ 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 $\mathbb{E}[\mathbf{u}_t\mathbf{u}_t'] = \boldsymbol{\Sigma} = (\mathbf{A}_0\mathbf{A}_0')^{-1}$. The matrices $\mathbf{B}$ and $\boldsymbol{\Sigma}$ are the reduced-form parameters, while $\mathbf{A}_0$ and $\mathbf{A}_+$ are the structural parameters. Similarly, $\mathbf{u}_t'$ are the reduced-form innovations, while $\boldsymbol{\varepsilon}_t'$ 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 parametrizations, we refer to Arias et al. (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.

2.1 Unconditional Forecasting

Assume that we want to forecast the observables for h periods ahead using the VAR in equations (2)-(4). Conditional on the history of data $\mathbf{y}^T = \left(\mathbf{y}_{1-p}' \dots \mathbf{y}_T'\right)'$, the forecast $\mathbf{y}_{T+1,T+h}' = \left(\mathbf{y}_{T+1}' \dots \mathbf{y}_{T+h}'\right)$ 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 absence of future shocks, whereas $\boldsymbol{\varepsilon}_{T+1,T+h}'\mathbf{M}$ is a stochastic part, reflecting the shocks unfolding over the forecast horizon with $\mathbf{M}$ capturing the associated impulse-response 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. in this case the choice of $\mathbf{Q}$ is irrelevant.

The distribution of the unconditional forecast can be interpreted as adding to the deterministic component of the forecast, $\mathbf{b}_{T+1,T+h}$, structural shocks drawn from their unconditional distribution.

In other words, the distribution of $\boldsymbol{\varepsilon}_{T+1,T+h}$ compatible with (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.

3 Conditional Forecasting and Scenario Analysis

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 sets of incorporating conditional assumptions into a forecast. In particular, we first consider the exercise conditioning on path for a subset of the variables; secondly conditioning on a particular path for the structural shocks; finally, we introduce the concept of structural scenario analysis which combines restrictions on observables and shocks. Despite their different interpretations, the three variants can all be written as linear restrictions on the VAR forecasts. Therefore, we first describe the general framework for incorporating linear restrictions, and then we proceed to show how each type of forecasting exercise can be written as a special case of the former.

3.1 General framework for restricted forecasts

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

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

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 $\boldsymbol{\Omega}_f$ 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

[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 et al. (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.

with mean $\mathbf{C}\mathbf{b}_{T+1,T+h}$ and variance $\mathbf{D}\mathbf{D}'$, where $\mathbf{D} = \mathbf{C}\mathbf{M}'$.

Andersson et al. (2010) show that the distribution of the forecast $\mathbf{y}_{T+1,T+h}$, conditional on a set of linear restrictions given by eq. (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 $\mathbf{D}$. Going back to the general expression of the forecast from a VAR model eq. (5), it is then possible to show that in order to retrieve the distribution of the conditional forecasts (9) one is effectively adding to the deterministic component of the forecast, $\mathbf{b}_{T+1,T+h}$, structural shocks distributed as

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

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 restricted forecasts are associated with a distribution of the structural shocks over the forecast horizon that deviates from their unconditional distribution, the standard Normal given by Equation 7.[5]

We now show how different methods to construct conditional forecasts are special cases of the

[5] See Appendix B in Andersson et al. (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 $\boldsymbol{\Omega}_f = \mathbf{0}$.

general framework above.

3.2 Conditional-on-Observables Forecasting

The classical conditional forecasting exercises, such as those first introduced by Doan et al. (1986), focus on calculating the forecast of some variables conditional on an exogenously imposed path for a subset of endogenous variables. We will call this environment 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 $\overline{\mathbf{C}}\mathbf{y}_{T+1,T+h} = \overline{\mathbf{f}}_{T+1,T+h}$ (cf. Waggoner and Zha (1999)) or more generally as density restrictions $\overline{\mathbf{C}}\mathbf{y}_{T+1,T+h} \sim \mathcal{N}\left(\overline{\mathbf{f}}_{T+1,T+h}, \overline{\mathbf{\Omega}}_f\right)$ (cf. Andersson et al. (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 $\mathbf{\Omega}_f = \overline{\mathbf{\Omega}}_f$. The solution for the conditional forecast of the entire set of variables is then given by Equation (9).

Note that, since the selection matrix $\overline{\mathbf{C}}$ does not depend of the parameters of the model the distribution of the conditional-on-observables forecasting only depends on the reduced-form parameters $\mathbf{B}$ and $\mathbf{\Sigma}$. Hence, given a set of reduced-form parameters, the choice of structural parameters, i.e. $\mathbf{Q}$ is irrelevant for retrieving the distribution of the conditional forecast (cf. Waggoner and Zha (1999), Proposition 1). Yet it is still 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 an identified structural model is not necessary to compute forecasts conditional-on-observables, if a structural model is available one can back out the values of the structural shocks implied by the conditional forecast.

3.3 Conditional-on-Shocks Forecasting

In this section we aim at constructing forecasts for all the endogenous variables of the system conditioning on a particular path of the structural shocks over the forecast horizon. We will call this environment conditional-on-shocks forecast.[6] For instance, Baumeister and Kilian (2014) use a

[6]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.

SVAR of the oil market to analyze the impact of a hypothetical oil supply shock. This practice is also used commonly used to produce conditional forecast with DSGE models (see Del Negro and Schorfheide (2013)).

Formally, let $\mathbf{\Xi}$ 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

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

where the $1 \times k_s$ vector $\mathbf{g}'_{T+1,T+h}$ and the conformable matrix $\mathbf{\Omega}_g$ denote the pre-specified mean restriction to the forecast of $\boldsymbol{\varepsilon}'_{T+1,T+h}\mathbf{\Xi}'$ and the associated variance.[7] In a Structural VAR 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 $\underline{\mathbf{C}}\mathbf{y}_{T+1,T+h}$

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

where $\underline{\mathbf{C}} = \mathbf{\Xi}(\mathbf{M}')^{-1}$. The restricted forecast associated with (15) is given by (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 $\mathbf{\Omega}_f = \mathbf{\Omega}_g$). The crucial difference with respect to the conditional-on-variable forecasts is that the linear restrictions, $\underline{\mathbf{C}}$, now depend on the impulse-response coefficients associated with the future shocks. Since $\mathbf{M}$ depends on $\mathbf{A}_0$, so will $\underline{\mathbf{C}}$, which implies that the identification will affect the conditional-on-shocks forecast of $\mathbf{y}_{T+1,T+h}$. It is obvious that, in order to impose restrictions upon their path, structural shocks need to be identified. For this reason, and unlike the unconditional and conditional-on-variable forecasts, the conditional-on-shocks forecasting will depend on the structural parameters $\mathbf{A}_0$ and $\mathbf{A}_+$. Hence, given a set of reduced-form

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

parameters, the choice of structural parameters, i.e. $\mathbf{Q}$ is now relevant.

Moreover, it is worth noticing that, contrary to what is the case when restrictions were placed only the observables, when conditional-on-shocks forecasting is used, the structural shocks that are not part of the conditioning exercise retain their unconditional (standard Normal) distribution.[8]

3.4 Structural Scenario Analysis

Conditional-on-shocks forecasting has the disadvantage that, the shocks being 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 shocks to generate the desired impact on a particular variable, or iterate between the shocks and the observed variables until achieving that result. 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 variables, together with a restriction that only a subset of the shocks can deviate from their unconditional distribution. The conditional-on-observables method implied restrictions on all structural shocks. Here, only the shocks that are assumed to be drivers of the scenario are allowed to deviate from the standard normal distribution, whereas the rest of the structural shocks that are not explicitly part of the 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 scenario, and therefore whose distribution is going to be restricted to be the same as their unconditional one.[9] Using the notation of Section 3.2, the restriction on the observables is implemented by imposing that

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

While using the notation of Section 3.3, the restriction on the structural shocks is implemented by

[8] See Appendix A for the full derivation.

[9] 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.

imposing that $\boldsymbol{\Xi}\boldsymbol{\varepsilon}_{T+1,T+h}$ is distributed as follows

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

The latter, in turn, implies $\underline{\mathbf{C}}\mathbf{y}_{T+1,T+h} \sim \mathcal{N}\left(\underline{\mathbf{C}}\mathbf{b}_{T+1,T+h}, \mathbf{I}_{n-p}\right)$, where $\underline{\mathbf{C}} = \boldsymbol{\Xi}(\mathbf{M}')^{-1}$. Taking the two set 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

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

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

Given these restrictions, the distribution of the restricted forecast is defined by equation (9), and 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 scenario must retain their unconditional distribution.

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 deviated 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 represented by is from the unconditional distribution of 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.[10] 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 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)-(14).

While it is straightforward to compute the KL divergence between $\mathcal{N}_{SS}$ and $\mathcal{N}_U$ using Equation 18, it is difficult to grasp whether any value for the KL divergence is large or small. In other words, the KL divergence can be easily used to 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 the KL divergence, McCulloch (1989) proposes to calibrate the KL divergence from two generic distributions P and Q to the the KL divergence between two easily interpretable distributions. In particular, he suggests to calibrate the KL divergence to the distance between two Bernoulli distributions, one with probability q and the other with probability 0.5: $D_{\mathrm{KL}}(\mathrm{Bern}(0.5)\|\,\mathrm{Bern}(q)) = D_{\mathrm{KL}}(\mathcal{N}_U\|\mathcal{N}_{SS})$. In this way any value for the KL divergence is translated into a comparison between the flip of a fair and a biased coin.

A drawback of McCulloch (1989)'s KL calibration in our setting is that the probability q is

[10] 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.

not scale invariant. Specifically, it increases quickly to one as nh, the dimension of the scenario, increases.[11] To solve this problem, we propose to calibrate the KL divergence (18) using two Binomial distributions instead of two Bernoulli distributions. Let $\mathcal{B}(m, p)$ denote the Binomial distribution that run 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 scenario.[12] Specifically, we find the probability $q \in [0.5, 1]$ so that $D_{\text{KL}}(\mathcal{B}(nh, 0.5) \| \mathcal{B}(nh, q)) = D_{\text{KL}}(\mathcal{N}_U \| \mathcal{N}_{SS})$:

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

where $z = D_{\text{KL}}(\mathcal{N}_U \| \mathcal{N}_{SS})$. The interpretation of q remains in line with the original McCulloch (1989)'s calibration. For example, a value of D_{KL} translating into a $q = 0.501$ suggests that the distribution of the structural shocks under the scenario considered is not at all far from the unconditional distribution of the shocks. Therefore the scenario considered is quite realistic. Similarly, a value of $q = 0.99$, suggests that the scenario requires a substantial deviation for the structural shocks from their unconditional distribution, it is extreme and therefore quite unlikely.

4 Algorithms to Implement the 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 incompletely identified SVARs. This algorithm can also be used to implement conditional-on-variables and conditional-on-shock forecasts as special cases. We assume normal-inverse-Wishart priors over the reduced-form parameters, hence the posterior distribution is also normal-inverse-Wishart as it extremely easy to generate draws from it. We also focus on sign restrictions to simply both the notation and the exposition. Let $\mathbf{S}_j$ be a $s_j \times r$ matrix of full row rank, where $0 \leq s_j$ where the $\mathbf{S}_j$ will define the sign restrictions on the j^{th} structural shock for

[11] In fact it is easy to show from (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.

[12] To see that the Binomial distribution effectively act as a rescaling of the KL divergence it is worth recalling that McCulloch (1989)'s KL calibration to a Bernoulli is $q(z) = (1 + \sqrt{1 - e^{-2z}})/2$.

$1 \leq j \leq n$. In particular, we assume that if $(\mathbf{A}_0, \mathbf{A}_+)$ satisfy the sign 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 $\mathbf{I}_n$. The algorithms can be easily modified to implement also zero restrictions by using the methods described in Arias et al. (2016b), and narrative sign restrictions as in Antolin-Diaz and Rubio-Ramirez (2016).

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 a orthogonal matrix, $\mathbf{Q}^{(i)}$, does not satisfy the sign 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 sign 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* $y_{T+1,T+h}^{(i)}$ *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 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 sign restrictions, the results of the structural scenario analysis will be robust across the set of structural models that satisfy the restrictions. This attractive feature of sign restrictions will come at the cost of very wide confidence bands around the forecast. Most importantly, there is a risk to include 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 et al. (2016a) and Antolin-Diaz and Rubio-Ramirez (2016), who highlight the importance of incorporating all the available identification restrictions that are uncontroversial.

Unlike unconditional forecasting and conditional-on-observables forecasting, structural scenario analysis requires that the structural parameters are identified in an economically meaningful way. Otherwise, the results of the scenario will not be reasonable. For instance, in the applications with monetary policy shocks that we present below, we find that a strategy based exclusively on contemporaneous sign restrictions often leads to implausibly large elasticites of observable variables to a monetary policy shock, mirroring the results of Kilian and Murphy (2012).[13] In the examples that follow we will propose to use a combination of traditional sign and zero restrictions at various

[13] 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.

horizons, as in Arias et al. (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 forecast scenarios.

5 Examples

5.1 Monetary Policy Scenarios

Let us 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, conditoning on observables and structural scenario analysis. Identification of the structural parameters is only necessary for the latter, as they are invariant to the structural identification used. However, by identifying the structural parameters first, we will be able to understand and interpret the results in light of their implications for the structural shocks.

We identify three structural shocks: a monetary policy (MP) shock, an aggregate demand shock (AD), and an aggregate supply shock (AS). 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 Monetary Policy (MP) shock reduces output and inflation and increases the federal funds rate on impact, and is restricted to have zero long-run impact on the level of output. We also restrict the monetary policy shock to have equal, and negative, impact on the inflation rate and the nominal interest rate at a horizon of 32 quarters. This serves to rule out "self-defeating" monetary policy shocks, which increase inflation in the long run, and imposes a zero long-run restriction on the real interest rate. Second, a contractionary Aggregate Demand (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 Aggregate Supply (AS) shock is restricted to reduce real GDP and increase inflation and the nominal interest rate on impact. Following Blanchard and Quah (1989), our identification scheme implies that the AS shock is the only one with a permanent effect on the level of output.

Table 1: 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.

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:

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.*

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

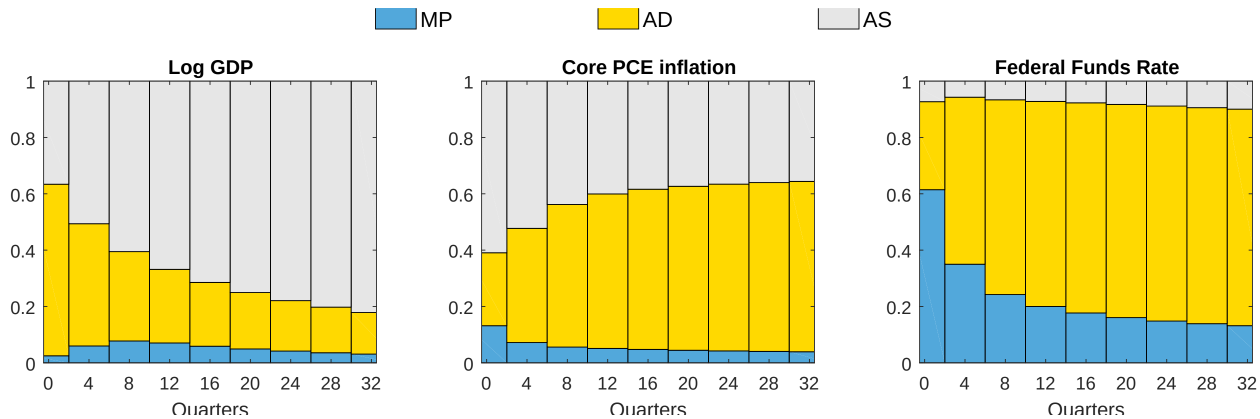

Note: The figure shows the mean posterior Forecast Error Variance Decomposition. For each panel, the colored bars represent the fraction of the total variance of the respective endogenous variable attributable to a specific structural shock at the horizon given by the horizontal axis.

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%, that inflation will recover gradually towards 2% and that the federal funds rate increases in a very gradual manner, approaching 2.5% 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 $t = T + 1 \dots T + 12$. The PDF of the unconditional distribution of the shocks, i.e., the standard normal distribution, is represented as a gray shaded area. As expected, it is clear from the PDFs that the unconditional forecast foresees that the future 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

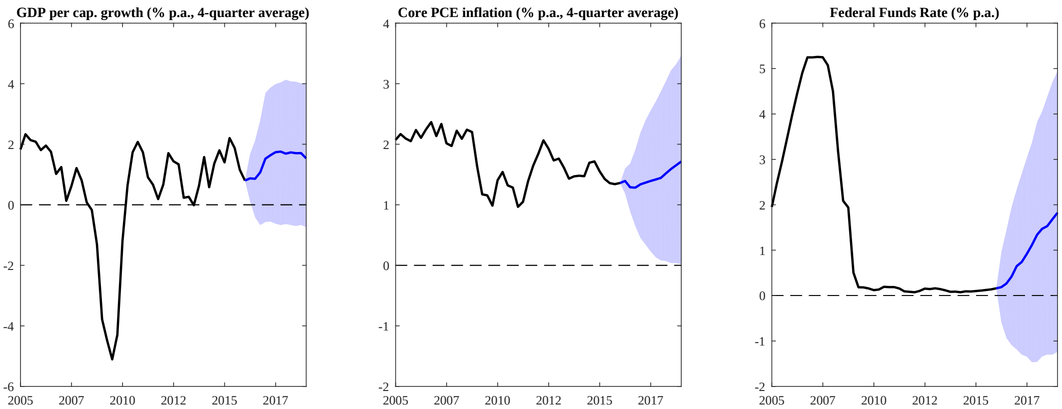

(b) Unconditional Structural Shocks

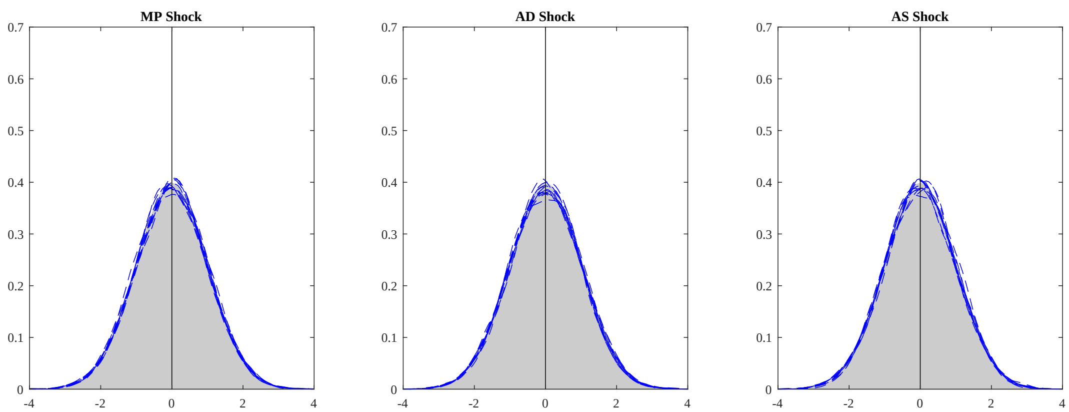

Note: In the top panel, for each column, the black solid lines represents actual data, the solid blue line is the median forecast and the blue shaded areas denote the 68 percent pointwise credible sets around the forecasts. The lower panel displays the PDFs of the structural shocks implied by the forecast 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 $t = T + 1 \dots T + 12$ the Federal Funds Rate increases by 50 basis points each quarter until it reaches 525 basis points. This pace of tightening of interest rate is identical to the one observed in the mid-2000s but substantially faster than the model's unconditional forecast (black dotted lines). Figure 4 considers the method of conditional-on-observables forecasting, as in Waggoner and Zha (1999) and section 3.2. [14] 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 $t = T + 1 \dots T + 12$ 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, given the sign restrictions imposed) 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 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.

[14] 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-Observable Forecast

(a) Conditional-on-Observable Forecasts

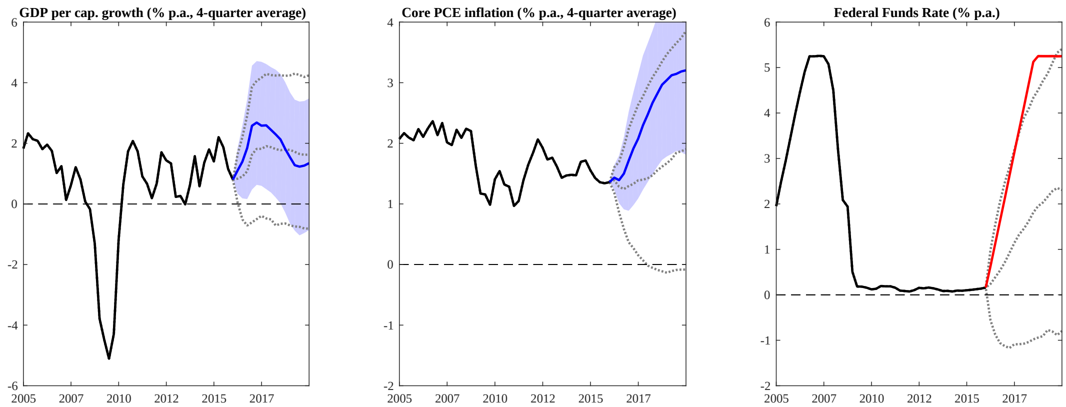

(b) Conditional-on-Observable Structural Shocks

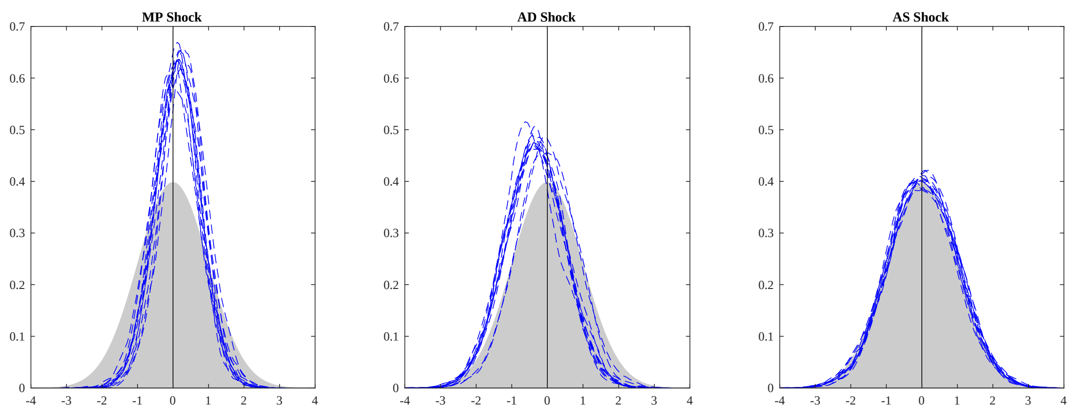

Note: In the top panel, for each column, the black solid lines represents actual data, the solid red line is the conditioning assumption on the observables, the solid blue line is the median forecast for the unrestricted variables and periods, and the blue shaded areas denote the 68 percent pointwise 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 $t = T+1 \dots T+12$. The gray shaded area is the PDF of a standard normal distribution.

Figure 5: Monetary Policy: Structural Scenario

(a) Structural Scenario Forecasts

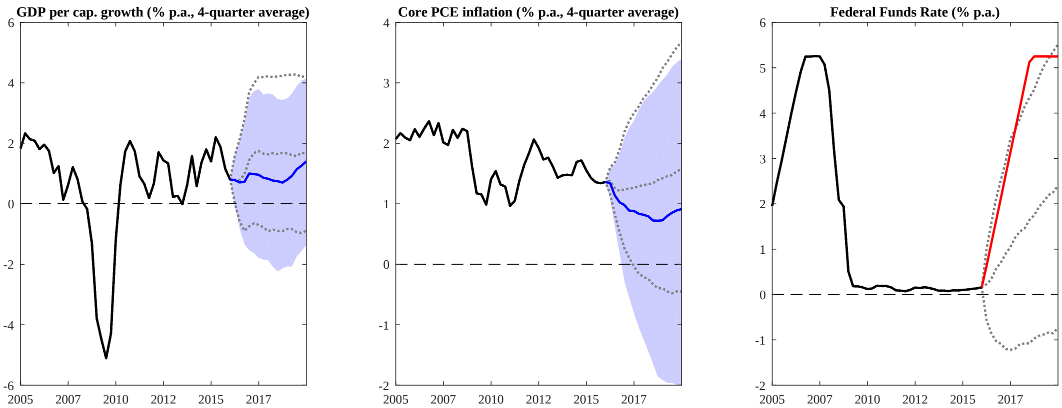

(b) Structural Scenario Shocks

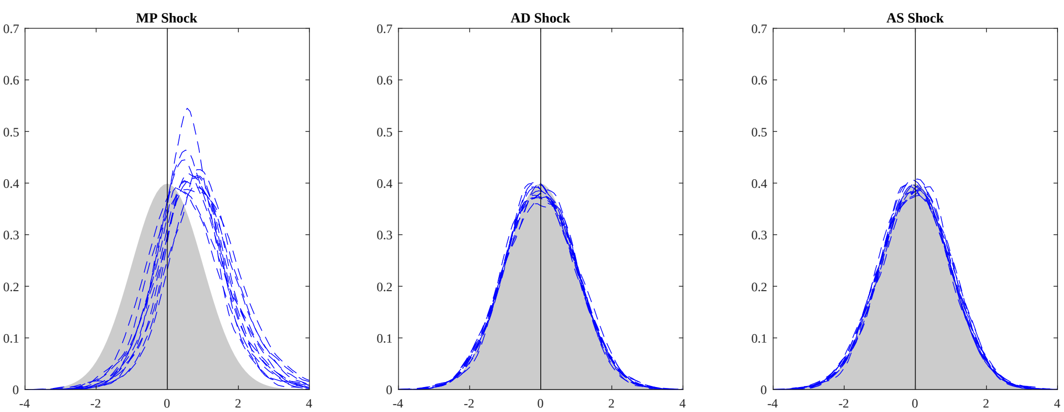

Note: In the top panel, for each column, the black solid lines represents actual data, the solid red line is the conditioning assumption on the observables, the solid blue line is the median forecast for the unrestricted variables and periods, and the gray shaded areas denote the 68 percent pointwise 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 $t = T+1 \dots T+12$. The gray shaded area is the PDF of a standard normal distribution.

5.1.3 Structural Scenario

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 that the monetary policy shock is the key driver of this scenario. In other words, the AD and AS supply 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% 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 policy shocks are required to produce the given path of the funds rate. These shocks exert a strong negative impact on output and inflation. As a result 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 on conditioning on a particular structural shock driving the scenario. 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% 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 estimation uncertainty of the reduced-form parameters. On the contrary, when the structural scenario is considered, the uncertainty about the structural parameters, in this case, the uncertainty about the precise impact of the monetary policy shock on output and inflation, must be taken into account. Our procedure fully incorporates the uncertainty both about the estimation of the reduced form parameters and about the structural model at hand.

5.1.4 Comparison of Policy Alternatives

Figure 6 analyzes some policy alternatives that might have been available to the monetary policy maker as of December 2015, when the Federal Reserve started increasing the Federal Funds rate. We consider three scenarios for the path of the interest rate and for all of them we make the scenario reflect different assumptions about the monetary policy shocks.[15] The "Baseline SEP" scenario is borrowed from the Summary of Economic Projections published by the FOMC.[16] It foresees the Federal Funds rate increasing by 25 basis points each quarter until it reaches 3.5 percent. The median forecast under this scenario is displayed as a red dashed line, and the 68 percent credible sets around the forecasts are plotted as red shaded areas. As can be seen from Panel (a), this scenario is associated with subdued inflation, which hovers just above 1 percent for the forecast horizon. Next, we consider an alternative scenario, which we denote "lower for longer", seen in Panel (b). In this 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, faster than in the baseline scenario. As can be seen from the figure, this 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" scenario of the previous subsection, in which the Fed Funds rate is increased 50 basis points every quarter until reaching 5.25 percent. As can be seen from Panel (c), this scenario is associated with low output and low inflation.

Table 2 looks at the plausibility of these alternative scenarios using the Kullback-Leibler divergence and its calibration as proposed in section 3.5. The first thing to notice is that when the 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, in the order of tens of thousands. The associated KL calibration is equal to 1 for all three cases, deeming the the 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 calibration confirm the intuition in Andersson et al. (2010) who suggest that conditioning on a

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

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

Figure 6: Comparison of Policy Alternatives (No uncertainty around the scenario)

(a) Baseline SEP

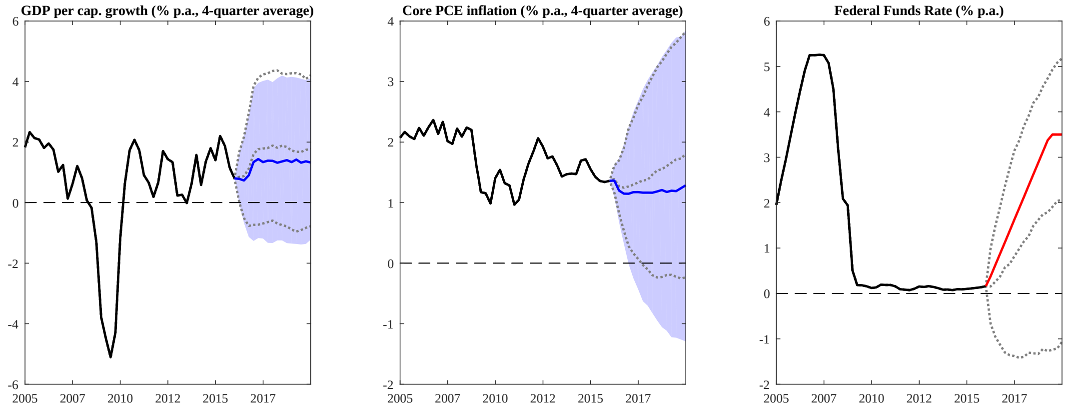

(b) Lower for Longer

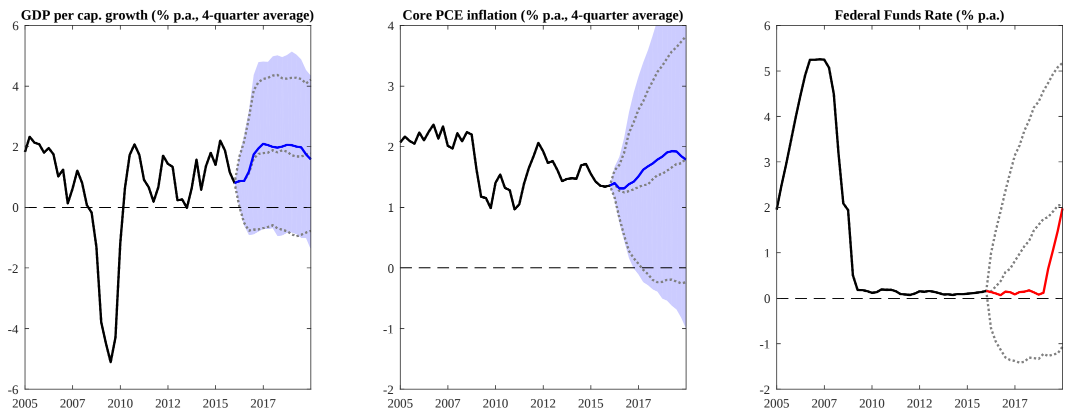

(c) Tighter

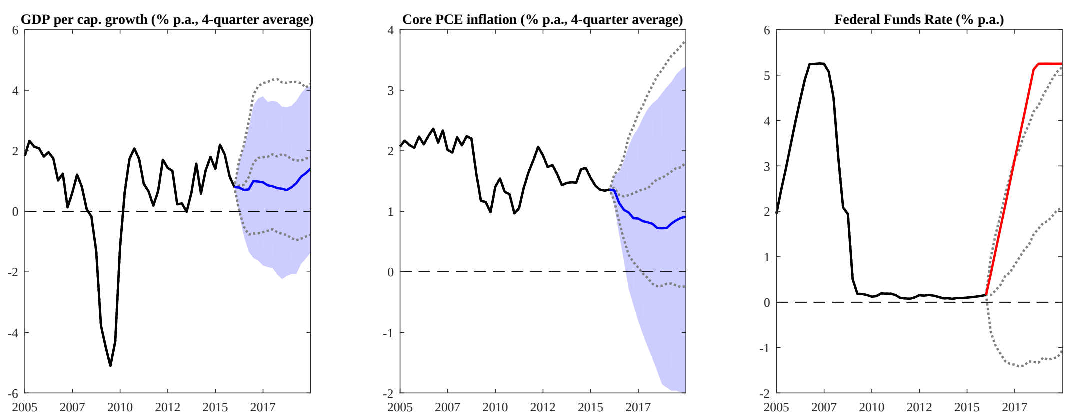

Table 2: Plausibility of the Alternative Scenarios

	Certain FFR		Uncertain FFR	
	KL divergence	Calibrated KL	KL divergence	Calibrated KL
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.

fixed path for a particular variable, ignoring the uncertainty about the conditioning assumption and can lead to unrealistic density forecasts. In Figure 7 we instead set the variance around the Fed Funds rate to its unconditional variance, $\boldsymbol{\Omega}_f = \overline{\mathbf{DD}}'$, as proposed by Andersson et al. (2010). In this case, we obtain a more realistic scenario, as seen in the last two columns of Table 2. Moreover, the uncertainty bands around the forecasts of the other variables are now wider. The values of the KL calibration further suggest that, whereas the baseline SEP forecast represent the most likely scenarios among the ones considered, the difference between the three is minimal and all of them are associated with important distortions in the distribution of the monetary policy shocks over the forecast horizon.

5.2 Stress-testing: the effect of a recession on asset prices and bank profitability

Our second example considers the impact of an economic recession on asset prices and bank profitability. Our key objective is to highlight that the potential impact of the same recessionary episode can be very different depending on the main driver of the recession. Specifically, a recession caused by financial shocks, like the one of 2007-09, can have more damaging impact on back profitability (and other financial variables) than non-financial recessions. 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 bill to 10-year government bond yield spread, the quarterly change

Figure 7: Comparison of Policy Alternatives (With uncertainty around the scenario)

(a) Baseline SEP

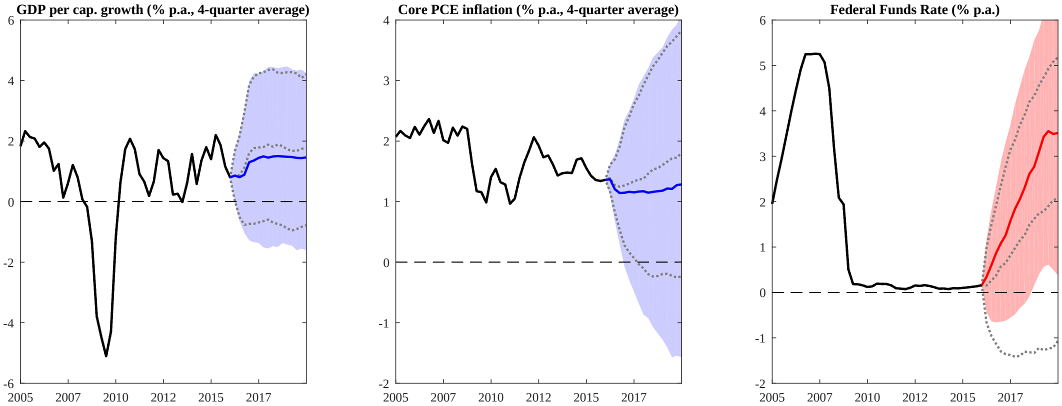

(b) Lower for Longer

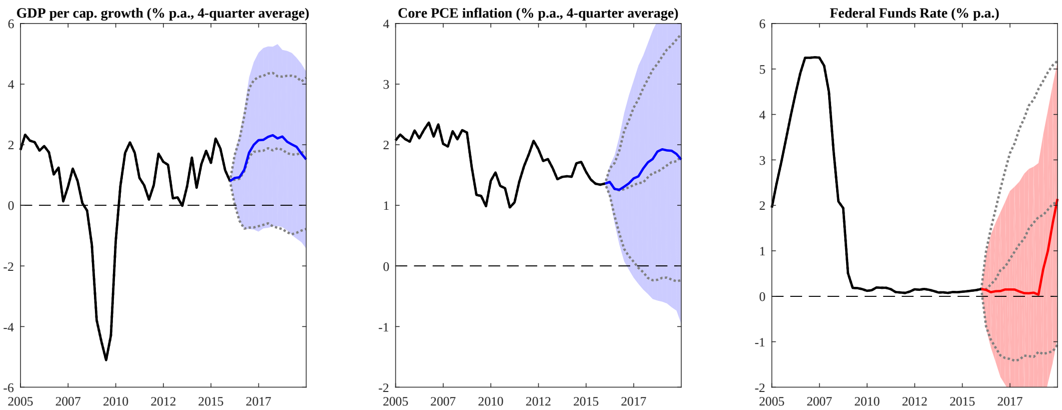

(c) Tighter

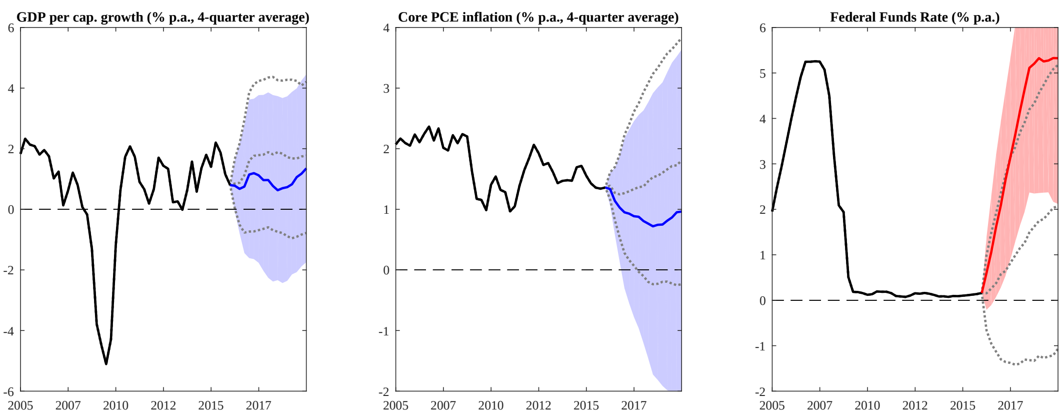

in the S&P 500 stock price index, the quarterly change in the S&P Case-Shiller House Price Index, the real price of oil, the BAA credit spread, the 3 month Treasury Bill-Eurodollar (TED) Spread; and an indicator of profitability in the banking system as a whole, the return on equity (ROE) of FDIC insured institutions.[17] The data is quarterly and seasonally adjusted, from 1984 to 2016.

As in the previous subsection, the identification scheme employs a combination of sign and narrative sign restrictions. The financial 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 of 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 in 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), and "directly touched off a run on money market funds" (p. 405), and triggered a large increase in spreads (p. 405). Following Bernanke's account 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

[17] The Haver mnemonics for the data are as follows: 3-month 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 is 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

Rule"[18] In this document, the Federal Reserve lays out various scenarios which banks participating in their stress tests can use to assess the impact 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%. The exact path of these two variables is given in Table 3.[19] 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. To implement this, we restrict the financial shocks to retain the $\mathcal{N}(0,1)$ distribution, and allow the other (unspecified) shocks to depart from their unconditional distribution. In other words, this scenario captures an economic downturn driven by shocks other than the financial shock. Figure 8 displays the results of the scenario. The dotted lines represent the unconditional forecast, and the solid blue lines and gray shaded areas are, respectively the median and 68% high posterior density intervals around the scenario. Inflation, interest rates, 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 scenario in which the financial shock is driving the recession, and the rest of the shocks retain the $\mathcal{N}(0,1)$ distribution. In this scenario of financial recession, 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 dramatically.

The exercise once again highlights the main point of this paper: when estimating what is the

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

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

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

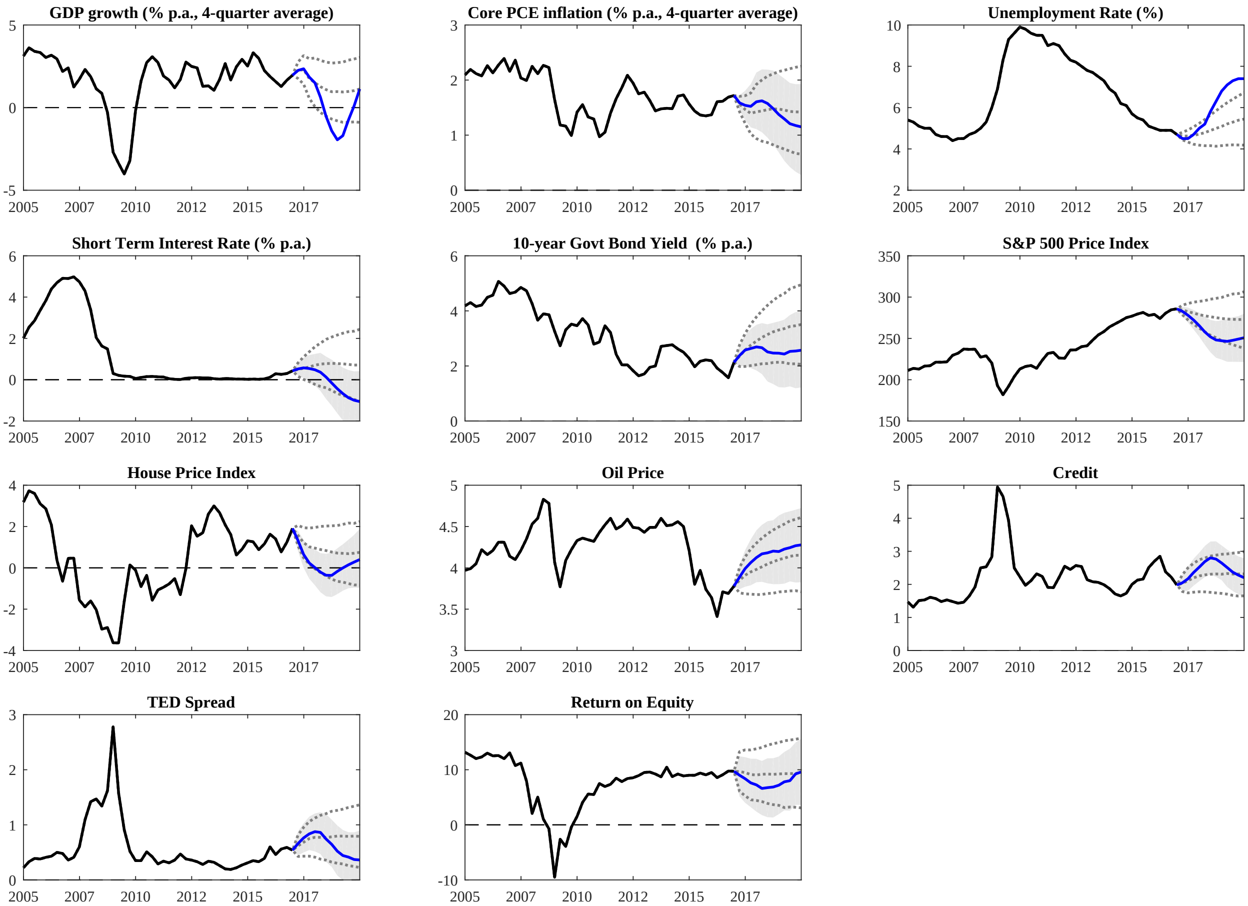

Note: For each panel, the dotted lines represent the unconditional forecast, and the solid blue lines and gray shaded areas are, respectively the median and 68% high posterior density intervals around the scenario.

implied path of some variables, given a pre-specified path for other variables, carefully considering which structural shock one has in mind can have important consequences for the answer. In the case of stress testing, the scenarios often try to estimate the impact of a recession on asset prices and bank profitability. But not all recessions are alike, and therefore what type of recession, i.e. what structural shock is driving the scenario, will matter.

Figure 9: Stress Test with Structural Scenario: Financial Recession

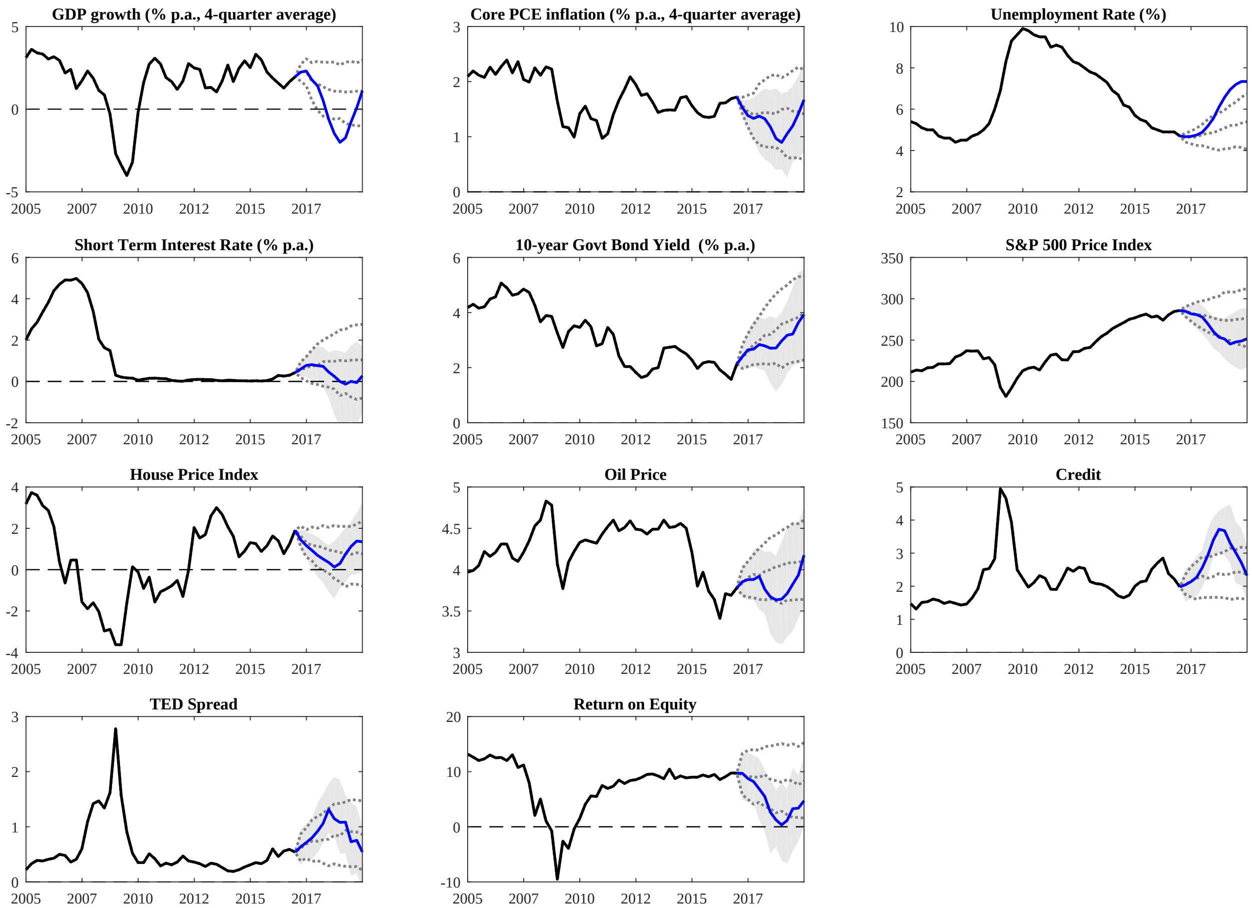

Note: For each panel, the dotted lines represent the unconditional forecast, and the solid blue lines and gray shaded areas are, respectively the median and 68% high posterior density intervals around the scenario.

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 endogenous 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, and assess the future value of the variables given some path for other variables which is driven by *a specific set of structural shocks.*

We have called this exercise structural scenario analysis.

A structural scenario has as point of departure a structural economic model. The literature has often undertaken this exercise through the specification of Dynamic Stochastic General Equilibrium models. Here we provide tools to perform structural scenario analysis using set and incompletely identified Structural VARs. This allows to fully capture uncertainty about the parameters and the structural model itself, while recent advances in identification allow to narrow down the set of structural models so as to produce meaningful economic results.

Given a structural VAR, the scenario analysis specifies a path for the future of some of the variables, and which structural shocks are the drivers of the scenario. In our illustrations, we considered the impact of alternative paths for the federal funds rate driven by exogenous monetary policy shocks, and stress-tested asset prices and bank profits to a recession caused by a financial shock. For the same path of the conditioning variables, the result can be very different depending on which shock is assumed to drive the scenario. Hence, structural scenario analysis can be a useful complement to conditional forecasting when an economic interpretation of the forecasts is sought.

A Appendix A: Vector notation for VAR forecasts

Assume that we want to forecast the observables for some period ahead using the VAR in eq. (**??**):

$$\mathbf{y}'_t = \sum_{\ell=1}^{p} \mathbf{y}'_{t-\ell}\mathbf{B}_\ell + \mathbf{c} + \mathbf{u}'_t \text{ for } 1 \leq t \leq T.$$

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

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

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

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

It is easy to see that, given that $\mathbf{B} = \mathbf{A}_+\mathbf{A}_0^{-1}$, and $\boldsymbol{\Sigma} = \left(\mathbf{A}_0\mathbf{A}'_0\right)^{-1}$, $\mathbf{M}'\mathbf{M}$ 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 $\mathbf{A}_+$.

Appendix 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 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-shock distribution of $\mathbf{y}_{T+1,T+h}$ is

$$\varepsilon_{T+1,T+h} \sim \mathcal{N}\left(\underline{\boldsymbol{\mu}}_{\varepsilon}, \underline{\boldsymbol{\Sigma}}_{\varepsilon}\right), \tag{20}$$

where

$$\underline{\boldsymbol{\mu}}_{\varepsilon} = \underline{\mathbf{D}}^*\left(\underline{\mathbf{C}}\mathbf{b}_{T+1,T+h} + \mathbf{g}_{T+1,T+h}\right) - \underline{\mathbf{D}}^*\underline{\mathbf{C}}\mathbf{b}_{T+1,T+h} = \underline{\mathbf{D}}^*\mathbf{g}_{T+1,T+h} \tag{21}$$

$$\underline{\boldsymbol{\Sigma}}_{\varepsilon} = \underline{\mathbf{D}}^*\boldsymbol{\Omega}_g\left(\underline{\mathbf{D}}^*\right)' + \underline{\hat{\mathbf{D}}}'\underline{\hat{\mathbf{D}}} \tag{22}$$

and $\underline{\mathbf{D}}^*$ is the generalized inverse of $\underline{\mathbf{D}} = \underline{\mathbf{C}}\mathbf{M}' = \boldsymbol{\Xi}$ and $\underline{\hat{\mathbf{D}}}$ 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 $\underline{\hat{\mathbf{D}}} = \left[\mathbf{0}_{k_s\times(nh-k_s)}, \mathbf{I}_{k_s}\right]'$. Using Equation (21), we have that the mean of the restricted structural shocks is

$$\underline{\boldsymbol{\mu}}_{\varepsilon} = \begin{bmatrix} \mathbf{g}_{t+1,t+h} \\ \mathbf{0}_{nh-k_s} \end{bmatrix}. \tag{23}$$

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 (22) simplifies to

$$\underline{\boldsymbol{\Sigma}}_{\varepsilon} = \begin{bmatrix} \boldsymbol{\Omega}_g & \mathbf{0}_{k_s \times (nh-k_s)} \\ \mathbf{0}_{(nh-k_s)\times k_s} & \mathbf{I}_{nh-k_s} \end{bmatrix}. \tag{24}$$

From Equations (23) and (24) it is easy to 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 its unconditional distribution.

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, C.E.P.R. Discussion Papers.

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," Working paper, Federal Reserve Bank of Atlanta.

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 & Compan.

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

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.

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

Del Negro, M. and F. Schorfheide (2013): *DSGE Model-Based Forecasting*, Elsevier, vol. 2 of *Handbook of Economic Forecasting*, chap. 0, 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.

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.

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.

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