

Manuscript version: Author's Accepted Manuscript

The version presented in WRAP is the author's accepted manuscript and may differ from the published version or Version of Record.

Persistent WRAP URL:

<http://wrap.warwick.ac.uk/136912>

How to cite:

Please refer to published version for the most recent bibliographic citation information. If a published version is known of, the repository item page linked to above, will contain details on accessing it.

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

© 2020 Elsevier. Licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International <http://creativecommons.org/licenses/by-nc-nd/4.0/>.



Publisher's statement:

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk.

Structural Scenario Analysis with SVARs¹

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

^a*London Business School*

^b*University of Warwick*

^c*Emory University, Federal Reserve Bank of Atlanta*

Abstract

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

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

1. Introduction

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

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

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

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

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

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

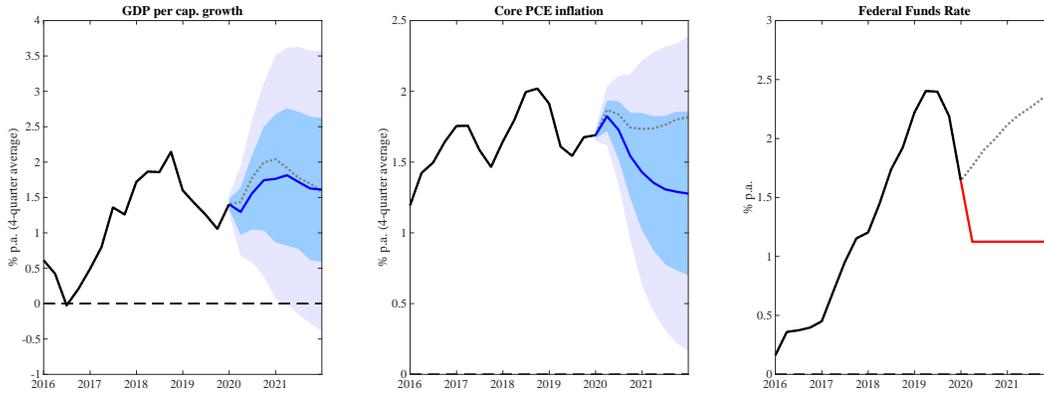
The differences are not surprising once we recognize that the bulk of the movements in the fed funds rate represents the systematic reaction of the Fed to output and inflation developments (see [Leeper et al., 1996](#)). The classic conditional forecasting question is really asking: “What is the most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

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

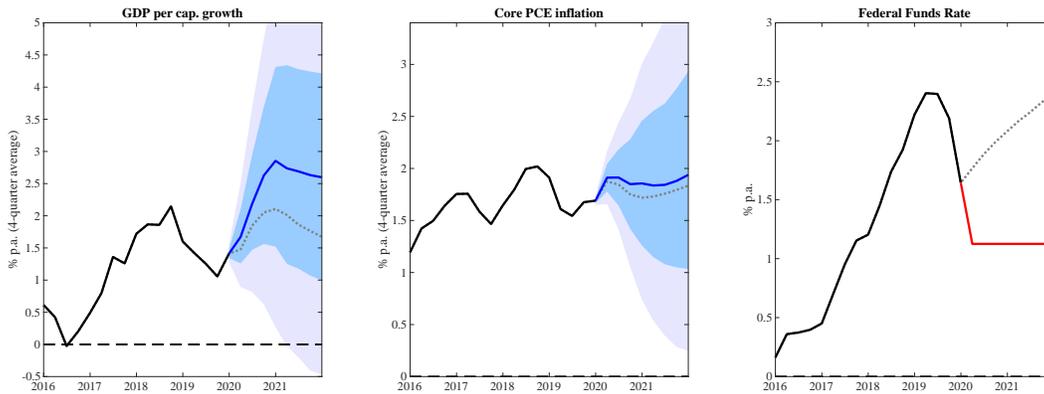
²Since 2009 the FOMC communicates interest rate targets as a range, e.g., “1-1.25%”. For brevity, we always refer to the lower end in the text, but we implement all our empirical scenarios using the mid-point.

Figure 1: ANSWERS TO TWO ALTERNATIVE QUESTIONS

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



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



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

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

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

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

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

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

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

2. Econometric framework

We work with the SVAR, written compactly

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

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

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

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

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

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

2.1. Unconditional forecasting

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

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

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

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

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

2.2. Conditional forecasts and structural scenarios: General framework

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

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

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

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

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

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

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

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

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

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

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

$$\boldsymbol{\Psi}_\varepsilon = \mathbf{D}^*\mathbf{\Omega}_f\mathbf{D}^{*'} - \mathbf{D}^*\mathbf{D}\mathbf{D}'\mathbf{D}^{*'}, \quad (11)$$

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

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

where \mathbf{D}^* is the Moore-Penrose inverse of \mathbf{D} . Because \mathbf{D} is full rank, the Moore-Penrose inverse always exists and is unique. From the definition of Ψ_ε , Equation (11) implies

$$\Sigma_\varepsilon = \mathbf{D}^* \Omega_f \mathbf{D}^{*'} + (\mathbf{I}_{nh} - \mathbf{D}^* \mathbf{D} \mathbf{D}' \mathbf{D}^{*'}), \quad (12)$$

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

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

$$\|\mathbf{C} \mathbf{b}_{T+1, T+h} + \mathbf{D} \boldsymbol{\mu}_\varepsilon - \mathbf{f}_{T+1, T+h}\| \quad \text{and} \quad \|\mathbf{D} (\mathbf{I}_{nh} + \Psi_\varepsilon) \mathbf{D}' - \Omega_f\|,$$

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Proof 1. See the online Appendix.

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

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

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

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

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

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

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

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

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

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

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

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

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

2.5. Special Case 3: Structural scenario analysis

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

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

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

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

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

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

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

3. How plausible is the structural scenario?

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

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

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

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

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

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

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

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

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

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

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

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

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

4. Algorithms for implementation

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

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

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

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

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

1. Conditioning on $\mathbf{y}^{T+h,(i-1)} = [\mathbf{y}^T, \tilde{\mathbf{y}}_{T+1,T+h}^{(i-1)}]$, use any valid algorithm to produce one draw from the structural parameters, $(\mathbf{A}_0^{(i)}, \mathbf{A}_+^{(i)})$ satisfying \mathbf{IR} .
2. Conditioning on $(\mathbf{A}_0^{(i)}, \mathbf{A}_+^{(i)})$ and \mathbf{y}^T , draw $\tilde{\mathbf{y}}_{T+1,T+h}^{(i)} \sim \mathcal{N}(\boldsymbol{\mu}_y, \boldsymbol{\Sigma}_y)$ satisfying \mathbf{SR} using Equations (13) and (14).
3. Return to Step 1 until the convergence and a sufficient number of draws have been obtained.

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

4.1. The importance of using all available identification restrictions

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

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

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

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

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

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

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

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

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

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

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

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

Consistent with the evidence that policy surprises are correlated across horizons, we aim to identify a shock that affects the level of the interest rate over the next year, without imposing zero impact of the anticipated shock on the contemporaneous fed funds rate.²³ We derive the restrictions above from the ones implied by introducing an equivalently defined anticipated monetary policy shock within the SW07 DSGE model.²⁴ The NSR improve the identification by bringing in information about announcements made by the Federal Reserve in 2003 and 2011-2012.²⁵ Notice that we only impose the sign of the shocks because our reading of these events indicates that contemporaneous news about growth and inflation may have also influenced the announcements. We wish to stress that in general the surprise variable \bar{S}_t is almost certainly affected in any quarter by many other shocks. Therefore, it would not constitute a valid instrument for anticipated monetary policy shocks. The advantage of the NSR approach in this context is that it allows us to incorporate the information that these shocks were affecting expectations during specific events.

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

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

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

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

Table 1: SUMMARY OF FORWARD GUIDANCE FOMC STATEMENT LANGUAGE

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

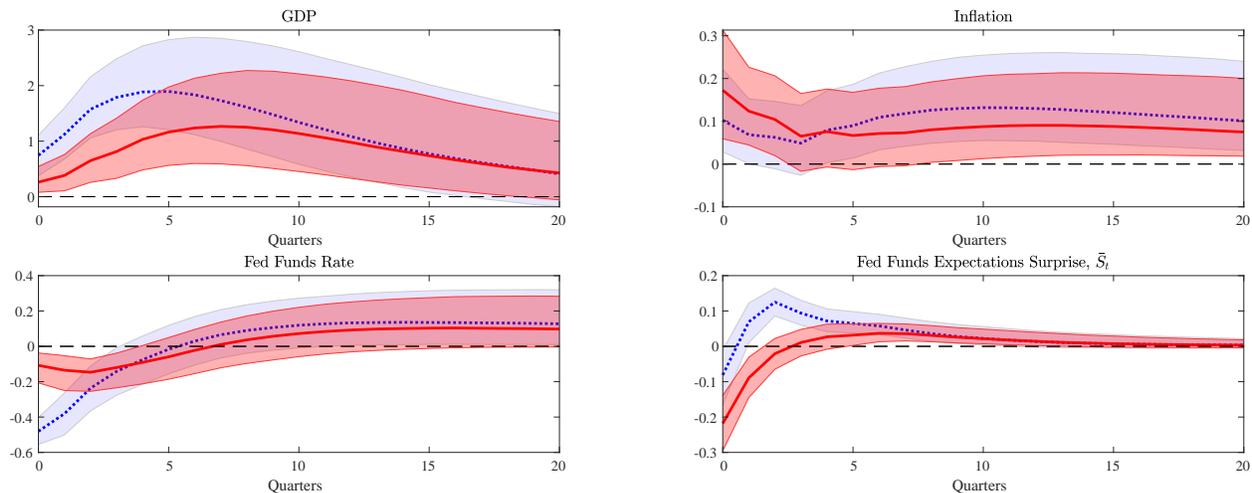
Identification Restriction 3: (*Aggregate Demand Shocks*) The aggregate demand shock has a positive effect on impact on output, consumption, investment, hours, wage, inflation and the fed funds rate, and zero long-run impact on the level of GDP.

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

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

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

Figure 2: MONETARY POLICY: IMPULSE RESPONSE FUNCTIONS TO ANTICIPATED VS UNANTICIPATED SHOCKS



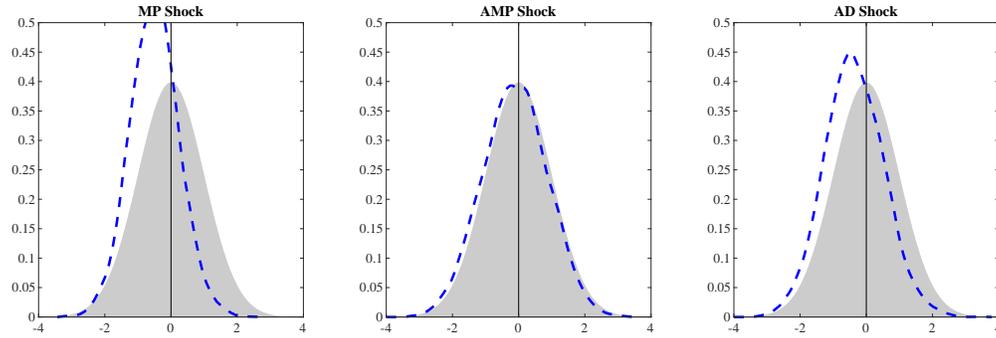
Note: The blue (lighter) shaded area represents the 68% (point-wise) HPD credible sets for the IRFs and the blue dotted lines are the median IRFs to an expansionary unanticipated monetary policy shock of one standard deviation. The red (darker) shaded areas and red solid lines display the equivalent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

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

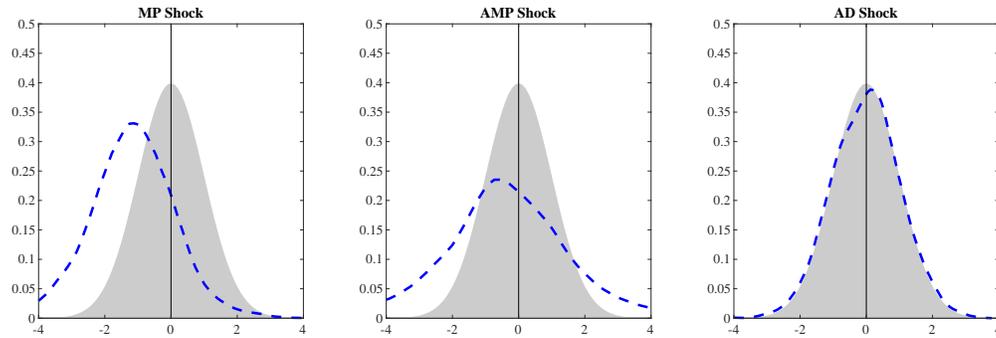
“[the] new language was meant to clarify that the maintenance of a very low fed

Figure 3: DISTRIBUTION OF SHOCKS UNDERLYING DIFFERENT EXERCISES

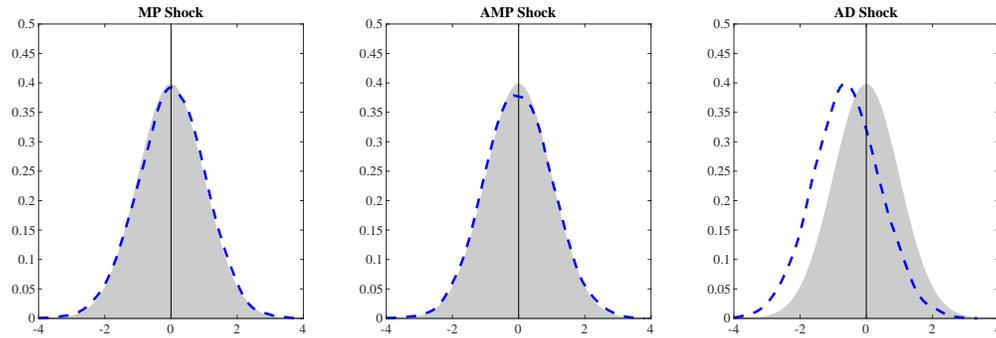
(a) Conditional-on-Observables Forecast



(b) An Odyssean Scenario



(c) A Delphic Scenario



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

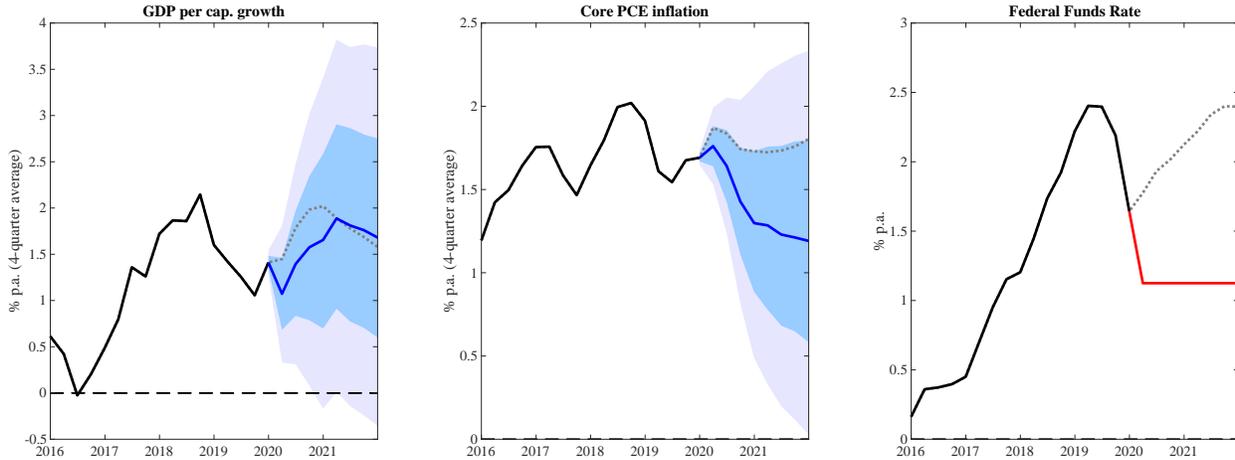
funds rate over that period did not reflect an expectation that the economy would remain weak, but rather reflected the Committee’s intention to support a stronger economic recovery.”²⁷

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

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

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

Figure 4: A DELPHIC SCENARIO



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

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

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

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

Table 2: RISK ASSESSMENT OF STRUCTURAL SCENARIOS

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

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

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

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

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

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

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

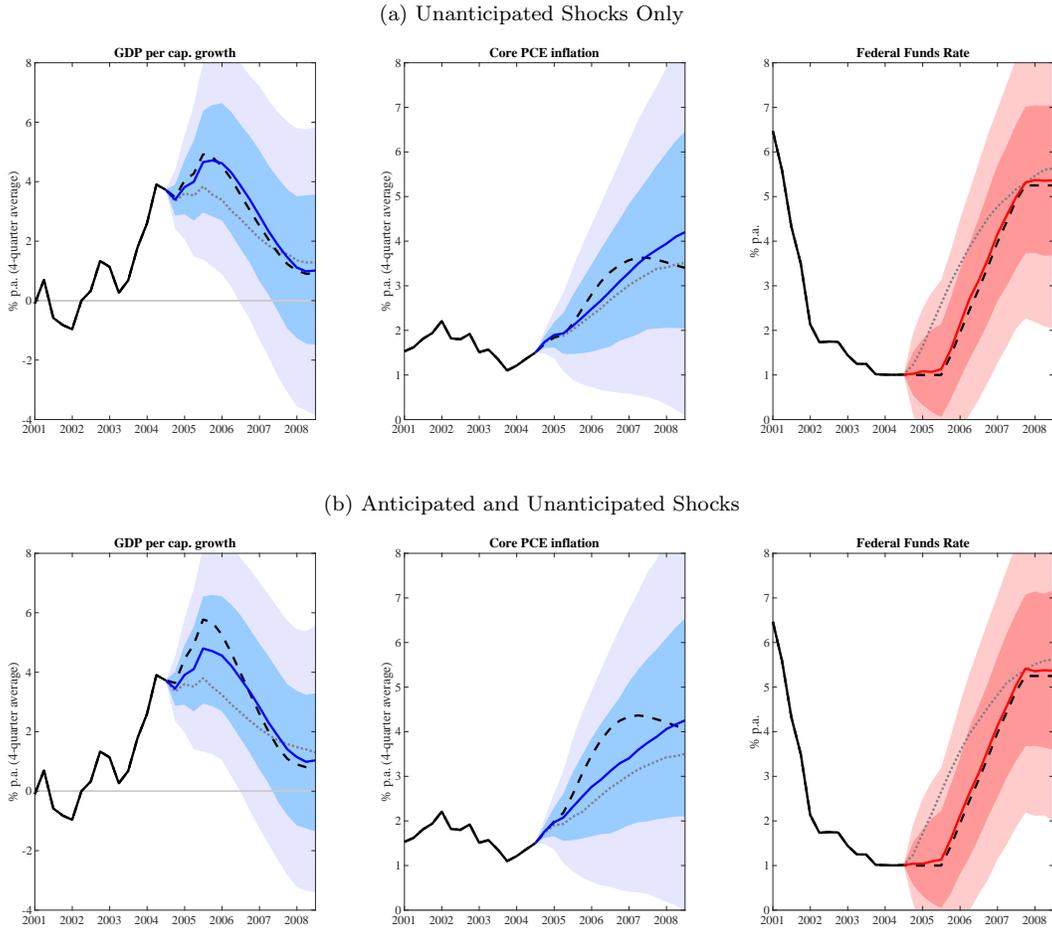
³¹The plausibility metric is 0.84 in the case where anticipation is allowed, and 0.8 when it is not, indicating that such exercises already imply a meaningful divergence from the unconditional distribution.

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

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

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

Figure 5: COMPARISON OF SVAR AND DSGE RESULTS



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

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

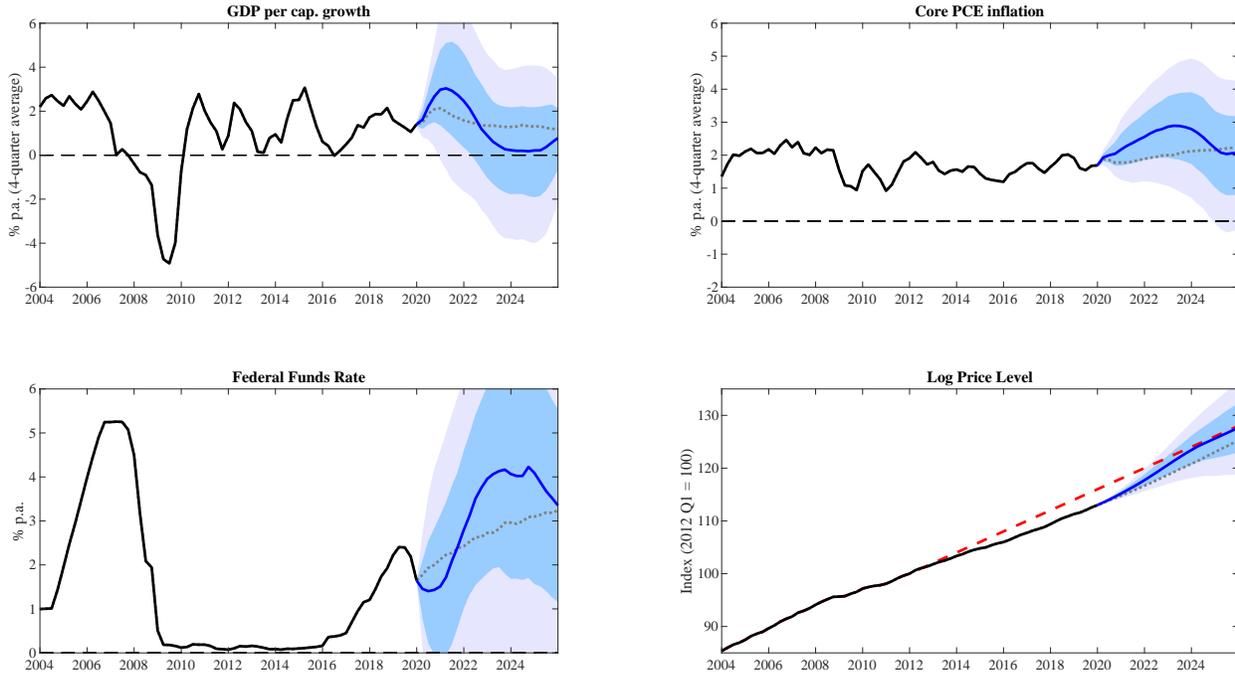
and inflation triggered by the new policy and bring inflation back to target. Using a car analogy, to catch up a car ahead of you that goes as fast as you, one needs to accelerate first and then decelerate. Decelerating the pace of price increases, requires higher interest rates to weaken aggregate demand.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

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

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

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

Figure 6: RETURNING TO THE PRICE LEVEL TARGET



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

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

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

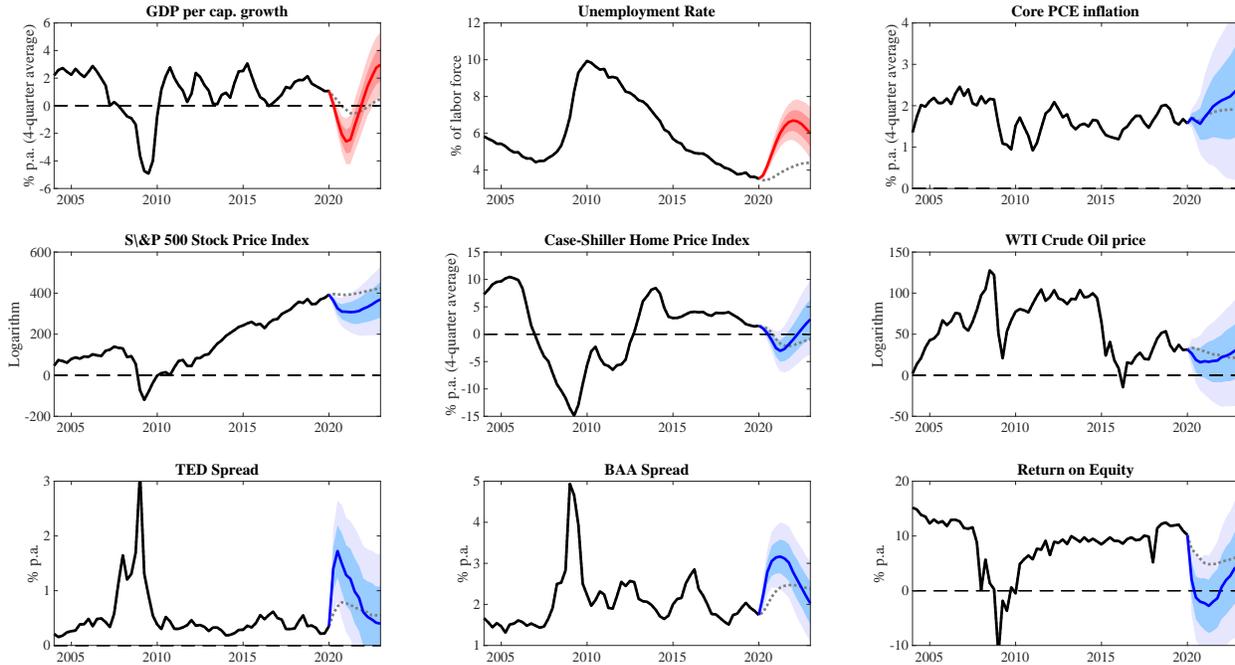
³⁴See the online Appendix for the definition of the variables and details of the model’s specification.

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

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

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

Figure 7: A STRESS TEST DRIVEN BY A FINANCIAL SHOCK



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

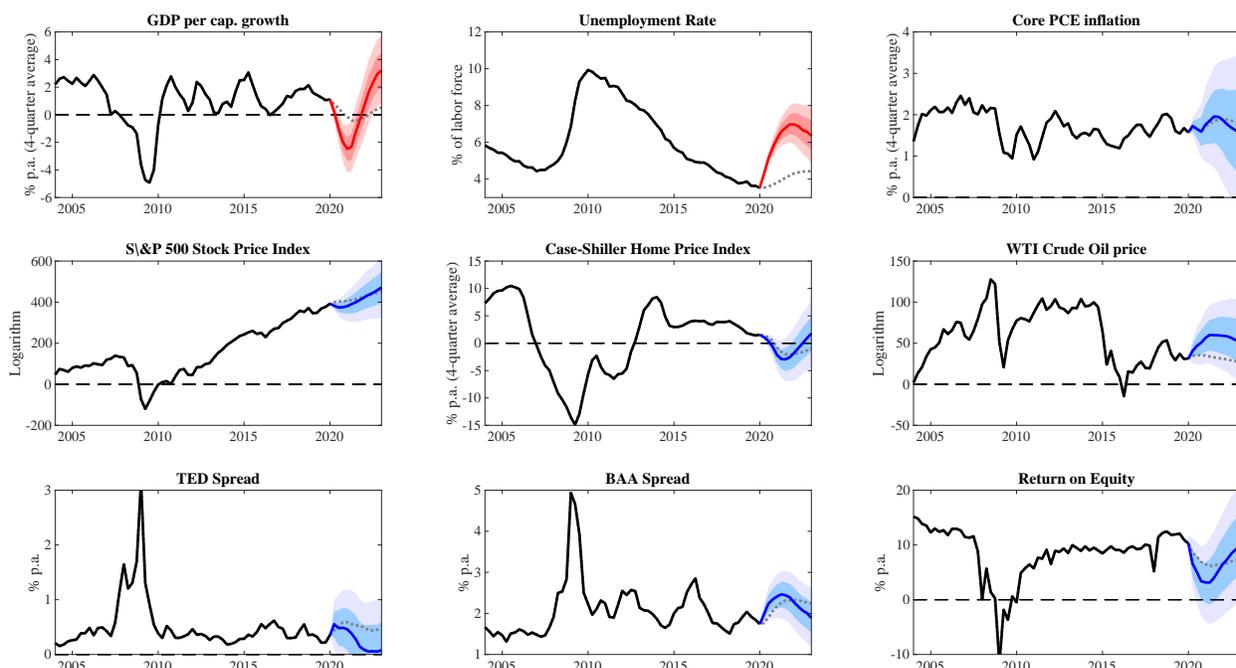
7. Conclusion

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

References

AASTVEIT, K. A., A. CARRIERO, T. E. CLARK, AND M. MARCELLINO (2017): “Have standard VARs remained stable since the crisis?” *Journal of Applied Econometrics*, 32, 931–951.

Figure 8: A STRESS TEST DRIVEN BY NON-FINANCIAL SHOCKS



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

AHMADI, P. A. AND H. UHLIG (2015): “Sign Restrictions in Bayesian FaVARs with an Application to Monetary Policy Shocks,” NBER Working Papers 21738, National Bureau of Economic Research, Inc.

ALTAVILLA, C., D. GIANNONE, AND M. LENZA (2016): “The Financial and Macroeconomic Effects of the OMT Announcements,” *International Journal of Central Banking*, 12, 29–57.

ANDERSSON, M. K., S. PALMQVIST, AND D. F. WAGGONER (2010): “Density Conditional Forecasts in Dynamic Multivariate Models,” Sveriges Riksbank Working Paper Series 247, Sveriges Riksbank.

ANTOLIN-DIAZ, J. AND J. F. RUBIO-RAMIREZ (2018): “Narrative Sign Restrictions for SVARs,” *American Economic Review*, 108, 2802–29.

ARIAS, J. E., D. CALDARA, AND J. F. RUBIO-RAMIREZ (2016): “The Systematic Component of Monetary Policy in SVARs: An Agnostic Identification Procedure,” International Finance Discussion Papers 1131, Board of Governors of the Federal Reserve System (U.S.).

ARIAS, J. E., J. F. RUBIO-RAMIREZ, AND D. F. WAGGONER (2018): “Inference Based on SVARs Identified with Sign and Zero Restrictions: Theory and Applications,” *Econometrica*, 86, 685–720.

BANBURA, M., D. GIANNONE, AND M. LENZA (2015): “Conditional forecasts and scenario analysis with vector autoregressions for large cross-sections,” *International Journal of Forecasting*, 31, 739–756.

- BAUMEISTER, C. AND J. D. HAMILTON (2015): “Sign restrictions, structural vector autoregressions, and useful prior information,” *Econometrica*, 83, 1963–1999.
- BAUMEISTER, C. AND L. KILIAN (2014): “Real-Time Analysis of Oil Price Risks Using Forecast Scenarios,” *IMF Economic Review*, 62, 119–145.
- BEN-ISRAEL, A. AND T. N. E. GREVILLE (2001): *Generalized Inverses: Theory and Applications*, Springer-Verlag New York, second ed.
- BERNANKE, B. S. (2015): *The Courage to Act*, W.W. Norton & Company.
- (2017): “Monetary Policy in a New Era,” Tech. rep., Peterson Institute for International Economics.
- BERNANKE, B. S., M. GERTLER, AND M. WATSON (1997): “Systematic Monetary Policy and the Effects of Oil Price Shocks,” *Brookings Papers on Economic Activity*, 28, 91–157.
- CAMPBELL, J. R., C. L. EVANS, J. D. FISHER, AND A. JUSTINIANO (2012): “Macroeconomic Effects of Federal Reserve Forward Guidance,” *Brookings Papers on Economic Activity*, 43, 1–80.
- CAMPBELL, J. R., J. D. FISHER, A. JUSTINIANO, AND L. MELOSI (2017): “Forward guidance and macroeconomic outcomes since the financial crisis,” *NBER Macroeconomics Annual*, 31, 283–357.
- CHRISTIANO, L. J., M. EICHENBAUM, AND C. L. EVANS (1999): “Monetary policy shocks: What have we learned and to what end?” in *Handbook of Macroeconomics*, ed. by J. B. Taylor and M. Woodford, Elsevier, vol. 1 of *Handbook of Macroeconomics*, chap. 2, 65–148.
- CLARIDA, R. H. (2019): “The Federal Reserve’s Review of Its Monetary Policy Strategy, Tools, and Communication Practices,” Speech, At “A Hot Economy: Sustainability and Trade-Offs,” a Fed Listens event sponsored by the Federal Reserve Bank of San Francisco, San Francisco, California.
- CLARIDA, R. H. AND D. COYLE (1984): “Conditional Projection by Means of Kalman Filtering,” NBER Technical Working Papers 0036, National Bureau of Economic Research, Inc.
- CLARK, T. E. AND M. W. MCCracken (2014): “Evaluating Conditional Forecasts from Vector Autoregressions,” Working Paper 1413, Federal Reserve Bank of Cleveland.
- CONTI, A. M., A. NOBILI, F. M. SIGNORETTI, ET AL. (2018): “Bank capital constraints, lending supply and economic activity,” Tech. rep., Bank of Italy, Economic Research and International Relations Area.
- DEL NEGRO, M., M. P. GIANNONI, AND C. PATTERSON (2012): “The forward guidance puzzle,” *FRB of New York Staff Report*.
- DEL NEGRO, M. AND F. SCHORFHEIDE (2013): *DSGE Model-Based Forecasting*, Elsevier, vol. 2A of *Handbook of Economic Forecasting*, chap. 2, 57–140.
- DOAN, T., R. B. LITTERMAN, AND C. A. SIMS (1986): “Forecasting and conditional projection using realistic prior distribution,” Staff Report 93, Federal Reserve Bank of Minneapolis.
- FERNÁNDEZ-VILLAVERDE, J., J. F. RUBIO-RAMÍREZ, T. J. SARGENT, AND M. W. WATSON (2007): “ABCs (and Ds) of Understanding VARs,” *American Economic Review*, 97, 1021–1026.
- GIACOMINI, R. AND G. RAGUSA (2014): “Theory-coherent forecasting,” *Journal of Econometrics*, 182, 145–155.
- GIANNONE, D., M. LENZA, D. MOMFERATOU, AND L. ONORANTE (2014): “Short-term inflation projections: a Bayesian vector autoregressive approach,” *International Journal of Forecasting*, 30, 635–644.

- GIANNONE, D., M. LENZA, H. PILL, AND L. REICHLIN (2012): “The ECB and the interbank market,” *The Economic Journal*, 122, F467–F486.
- GIANNONE, D., M. LENZA, AND L. REICHLIN (2019): “Money, Credit, Monetary Policy, and the Business Cycle in the Euro Area: What Has Changed Since the Crisis?” *International Journal of Central Banking*, 15, 137–173.
- HAMILTON, J. D. AND A. M. HERRERA (2004): “Oil Shocks and Aggregate Macroeconomic Behavior: The Role of Monetary Policy: Comment,” *Journal of Money, Credit and Banking*, 36, 265–286.
- JAROCIŃSKI, M. (2010): “Conditional forecasts and uncertainty about forecast revisions in vector autoregressions,” *Economics Letters*, 108, 257–259.
- KILIAN, L. AND L. T. LEWIS (2011): “Does the Fed Respond to Oil Price Shocks?” *Economic Journal*, 121, 1047–1072.
- KILIAN, L. AND S. MANGANELLI (2008): “The central banker as a risk manager: Estimating the Federal Reserve’s preferences under Greenspan,” *Journal of Money, Credit and Banking*, 40, 1103–1129.
- KILIAN, L. AND D. P. MURPHY (2012): “Why agnostic sign restrictions are not enough: Understanding the dynamics of oil market VAR models,” *Journal of the European Economic Association*, 10, 1166–1188.
- LEEPER, E. M., C. A. SIMS, AND T. ZHA (1996): “What does monetary policy do?” *Brookings Papers on Economic Activity*, 1996, 1–78.
- LEEPER, E. M. AND T. ZHA (2003): “Modest policy interventions,” *Journal of Monetary Economics*, 50, 1673–1700.
- LUCAS, R. E. (1976): “Econometric policy evaluation: A critique,” in *Carnegie-Rochester Conference Series on Public Policy*, vol. 1, 19–46.
- LUDVIGSON, S. C., S. MA, AND S. NG (2017): “Shock Restricted Structural Vector-Autoregressions,” NBER Working Papers 23225, National Bureau of Economic Research, Inc.
- MCCULLOCH, R. E. (1989): “Local Model Influence,” *Journal of the American Statistical Association*, 84, 473–478.
- MCKAY, A., E. NAKAMURA, AND J. STEINSSON (2016): “The power of forward guidance revisited,” *American Economic Review*, 106, 3133–58.
- PENROSE, R. (1955): “A generalized inverse for matrices,” *Mathematical Proceedings of the Cambridge Philosophical Society*, 51, 406–413.
- (1956): “On best approximate solutions of linear matrix equations,” *Mathematical Proceedings of the Cambridge Philosophical Society*, 52, 17–19.
- RAMEY, V. A. (2016): “Macroeconomic Shocks and Their Propagation,” NBER Working Papers 21978, National Bureau of Economic Research, Inc.
- ROBERTSON, J. C., E. W. TALLMAN, AND C. H. WHITEMAN (2005): “Forecasting Using Relative Entropy,” *Journal of Money, Credit and Banking*, 37, 383–401.
- RUBIO-RAMIREZ, J. F., D. F. WAGGONER, AND T. ZHA (2010): “Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference,” *Review of Economic Studies*, 77, 665–696.
- SIMS, C. A. (1986): “Are forecasting models usable for policy analysis?” *Quarterly Review*, 2–16.

- (1993): “A Nine-Variable Probabilistic Macroeconomic Forecasting Model,” in *Business Cycles, Indicators and Forecasting*, National Bureau of Economic Research, Inc, NBER Chapters, 179–212.
- SMETS, F. AND R. WOUTERS (2007): “Shocks and frictions in US business cycles: A Bayesian DSGE approach,” *American Economic Review*, 97, 586–606.
- TALLMAN, E. W. AND S. ZAMAN (2018): “Combining Survey Long-Run Forecasts and Nowcasts with BVAR Forecasts Using Relative Entropy,” *Federal Reserve Bank of Cleveland Working Papers*.
- WAGGONER, D. F. AND T. ZHA (1999): “Conditional Forecasts in Dynamic Multivariate Models,” *Review of Economics and Statistics*, 81, 639–651.