The analytics of the sign restriction approach to shock identification: a framework for understanding the empirical macro puzzles*

Dario Caldara† Christophe Kamps‡

February 15, 2010

Abstract

The empirical literature using vector autoregressive (VAR) models has uncovered a number of important macro "puzzles", such as the price puzzle. This paper shows that the sign-restriction approach to shock identification is well-suited for understanding when and why the puzzles arise. We show that for bivariate VAR models it is possible to analytically characterize the entire set of pure sign restriction solutions and to divide it into a subset of "no-puzzle" solutions and a subset of "puzzle" solutions. We establish a simple condition for non-emptiness of the subset of "puzzle" solutions. We also show how this condition can be generalized to multivariate VAR models with multiple sign restrictions. In this context, we establish the "seemingly unrestricted variable" (SUV) property of sign restrictions, showing that those variables whose responses to a given shock are not explicitly sign restricted will be implicitly sign restricted under quite general conditions. We further show that the standard numerical penalty function used in the literature to select one out of the infinitely many admissible pure sign restriction solutions is nested in a general objective function. We show analytically that the optimum of the standard numerical penalty function is a "no-puzzle" solution.

We illustrate the relevance of our methodological findings with the example of fiscal policy shock identification. In the literature there is considerable disagreement on the size and even sign of fiscal multipliers. Our framework allows understanding the sources of disagreement as we are able to analytically characterize the set of fiscal multipliers consistent with standard identifying assumptions used in the literature to identify fiscal policy shocks.

JEL classification: C32; C52; E62.

Keywords: vector autoregression; identification; sign restrictions; price puzzle; fiscal policy.

*Preliminary and incomplete. Please do not quote without permission. The views expressed in this paper do not necessarily reflect those of the European Central Bank.

†Institute for International Economic Studies, Stockholm University, Stockholm, Sweden. Email: dario.caldara@iies.su.se.

‡European Central Bank, Frankfurt am Main, Germany. Email: christophe.kamps@ecb.int.

1 Introduction

Following the seminal contribution by Sims (1980), structural VAR models have become a standard tool to assess the dynamic macroeconomic effects of shocks and to cross-check the implications of dynamic general equilibrium models. Over the years, the empirical literature using VAR models has established a number of uncontroversial stylized facts but at the same time has uncovered a number of important macro "puzzles" such as the price puzzle (Sims (1992)), the output puzzle (Uhlig (2005)), the liquidity effect puzzle (Leeper and Gordon (1992)), the delayed overshooting puzzle (Eichenbaum and Evans (1995)), the exchange rate puzzle (Grilli and Roubini (1996)), the government-private-consumption puzzle (Galí, López-Salido, and Vallés (2007)) and the tax multiplier puzzle (Perotti (2005)).

To be sure, whether an empirical result is viewed as puzzling depends on the observer's theoretical prior. What looks counterintuitive to one observer might make perfect sense for another observer. Also, it is not clear whether a puzzling empirical result should be viewed as a challenge for existing theory or rather as a reason to rethink the empirical identification strategy. Indeed, the robustness of the empirical macro puzzles has been and remains controversial. Yet, the fact that these puzzles appear frequently in empirical studies and that they are indeed quite numerous raises the question of the fundamental reasons explaining the existence of puzzling results. This paper shows that the sign-restriction approach to shock identification introduced by Uhlig (2005) is well-suited for understanding when and why the puzzles arise. What is sometimes viewed as one of the main weaknesses of the sign-restriction approach, namely that sign restrictions do not pin down a unique structural model, turns out to be a crucial advantage in this context.

The macro puzzles listed above share one important common feature: the puzzle consists in a counterintuitive response of a non-policy variable (output, prices, private consumption, exchange rate) to an unexpected (monetary or fiscal) policy shock. Against this background, we adopt the following modeling strategy. We start with the simplest conceivable framework: a bivariate vectorautoregression with two variables, a non-policy variable and a policy variable. We further assume that two fundamental shocks are driving the dynamics of these two variables, a non-policy shock and a policy shock. The object of interest is the response of the non-policy variable to the policy shock. In this setup the non-policy shock plays a crucial role because it controls for comovements in the two variables that are unrelated to unanticipated changes in policy but originate from non-policy shocks. Consider two prominent examples:

First, the macroeconomic effects of monetary policy shocks. Assume that the short-term interest rate is the policy instrument and inflation is the non-policy variable. Mainstream macroeconomics predicts that an unanticipated increase in interest rates leads to a decrease in inflation. Instead, empirical studies based on structural VAR models often find that inflation increases in response to a monetary policy contraction, giving rise to what is commonly referred to as the *price puzzle* (see e.g. Sims (1992) and Uhlig (2005)). Why and when

does the price puzzle arise? It can arise because empirically the correlation between interest rate and inflation residuals is positive. This positive correlation potentially has two sources: (i) the systematic response of monetary policy to non-policy shocks that lead to deviations of inflation from the policy target and (ii) the effect of contractionary monetary policy shocks on inflation.

To preview one of the main results to be established in this paper: the price puzzle arises if and only if the non-policy shock explains less than 100% of the covariance between interest rate and inflation residuals. In this case the policy shock has to explain some of the positive comovement, which is possible if and only if the interest rate and inflation move in the same direction in response to a monetary policy shock. In Section 3 we analytically characterize both the set of pure sign-restriction solutions exhibiting the "puzzle" and the set of pure sign-restriction solutions conforming with the mainstream view. Among other things, we establish a simple condition under which the set of "puzzle" solutions is empty. This condition is unlikely to be satisfied empirically, explaining why many attempts to resolve the price puzzle have not proved robust (Hanson (2004)).[1]

Second, the macroeconomic effects of fiscal policy shocks. Assume that taxes are the policy instrument[2] and output is the non-policy variable. Will an unanticipated tax cut stimulate output? Theory is inconclusive. While Keynesian and neoclassical theory predict output increases, Ricardian equivalence predicts output neutrality and non-Keynesian theory even predicts that output will be depressed under certain conditions. Available empirical evidence is inconclusive, too. While Blanchard and Perotti (2002) find that unanticipated tax cuts significantly and persistently stimulate output in the U.S., Perotti (2005) finds that in the short run tax cuts depress output in the U.S. and in the medium to long run they have insignificant effects, giving rise to the *tax multiplier puzzle*. Our analytical results show that the short-run tax cut multiplier will be positive if and only if the cyclical adjustment is sufficiently strong, i.e. if and only if the non-policy shock explains more than 100% of the covariance between tax and output residuals.[3]

More fundamentally, our paper makes the following methodological contributions. We show that for bivariate VAR models it is possible to analytically characterize the entire set of pure sign restriction solutions and to divide it into a subset of "no-puzzle" solutions and a subset of "puzzle" solutions. We establish a simple condition for non-emptiness of the subset of "puzzle" solutions. We also show how this condition can be generalized to multivariate VAR models with multiple sign restrictions. In this context, we establish the "seemingly unrestricted variable" (SUV) property of sign restrictions, showing that those

[1] One approach that has to be successful by definition is to impose a sign restriction on the inflation response to a monetary policy shock such as to rule out puzzle solutions (Uhlig (2005)).

[2] We take taxes as an example. Taxes and government spending are modelled jointly in Section 5.

[3] Our results also show that the subset of positive tax cut multipliers is sufficiently large to explain the strongly differing estimates of Blanchard and Perotti (2002) and Romer and Romer (2010).

variables whose responses to a given shock are not explicitly sign restricted will be implicitly sign restricted under quite general conditions. We further show that the standard numerical penalty function used in the literature to select one out of the infinitely many admissible pure sign restriction solutions is nested in a general objective function. We show analytically that the optimum of the standard numerical penalty function is a "no-puzzle" solution.

We illustrate the relevance of our methodological findings with the example of fiscal policy shock identification. In the literature there is considerable disagreement on the size and even sign of fiscal multipliers. Our framework allows understanding the sources of disagreement as we are able to analytically characterize the set of fiscal multipliers consistent with standard identifying assumptions used in the literature to identify fiscal policy shocks. We show that typical reduced-form VAR models and sign restrictions are consistent with both positive and negative tax cut multipliers, contradicting conjectures ruling out negative tax cut multipliers (Mountford and Uhlig (2009)). We further show that the identification of government spending shocks requires imposing parametric restrictions in addition to the standard sign restrictions. In the absence of such parametric restrictions the SUV property of sign restrictions leads to counterfactual cyclical properties of government spending, heavily distorting the size of government spending multipliers. We show that, controlling for this distortion, a prominent finding in the recent empirical literature, namely that deficit-financed tax cuts are more effective than deficit-spending in stimulating ouput (Mountford and Uhlig (2009)), vanishes.

The remainder of this paper is organized as follows. Section 2 introduces the reduced-form VAR model and the sign-restriction approach to shock identification. Section 3 establishes key results for bivariate VAR models. Section 4 shows under which conditions these results can be generalized to the multivariate case with multiple sign restrictions. Section 5 applies our methodological framework to the fiscal multiplier debate. Section 6 concludes.

2 The identification problem

2.1 The reduced-form VAR model

We consider the pth-order zero-mean vector autoregression

$$x_t = B_1 x_{t-1} + B_2 x_{t-2} ... + B_p x_{t-p} + u_t, \tag{1}$$

where x_t is a $(k \times 1)$ vector collecting the set of model variables and u_t is a $(k \times 1)$ vector of one-step-ahead prediction errors with $E[u_t] = 0$ and $(k \times k)$ symmetric covariance matrix $E[u_t u_t'] = \Sigma_u = [\sigma_{ij}]$. We denote the standard deviation of prediction error u_{it} as σ_i, with $\sigma_i \equiv \sqrt{\sigma_{ii}}$, $i = 1, ..., k$. We further denote the correlation coefficient between errors u_{it} and u_{jt} as ρ_{ij}, with $\rho_{ij} \equiv \frac{\sigma_{ij}}{\sigma_i \sigma_j}$, $i, j = 1, ..., k$.

Assumption (Covariance matrix is positive definite)**.** *The prediction error covariance matrix Σ_u is a real symmetric positive definite matrix.*

The positive definiteness of Σ_u implies $\sigma_i > 0$ for all i and $|\sigma_{ij}| < \sigma_i \sigma_j$ or, equivalently, $|\rho_{ij}| < 1$ for $i \neq j$. This assumption guarantees that the Cholesky factorization of the covariance matrix exists, i.e. a unique lower triangular matrix $\tilde{A} \in R^{k \times k}$ with positive elements on the principal diagonal exists such that $\Sigma_u = \tilde{A}\tilde{A}'$ (Golub and van Loan (1996), Theorem 4.2.5, p. 143).

We write the individual elements of the covariance matrix in terms of standard deviations and correlation coefficients, which will prove useful for the presentation of the analytical results in the remainder of the paper:

$$\underset{(k \times k)}{\Sigma_u} = \begin{bmatrix} \sigma_1^2 & \dots & \sigma_1\sigma_j\rho_{1j} & \dots & \sigma_1\sigma_k\rho_{1k} \\ \vdots & \ddots & \vdots & & \vdots \\ \sigma_1\sigma_j\rho_{1j} & \dots & \sigma_i^2 & \dots & \sigma_j\sigma_k\rho_{kj} \\ \vdots & & \vdots & \ddots & \vdots \\ \sigma_1\sigma_k\rho_{1k} & \dots & \sigma_j\sigma_k\rho_{kj} & \dots & \sigma_k^2 \end{bmatrix}. \tag{2}$$

2.2 The sign-restriction approach to identification

The identification problem consists in translating the correlated prediction errors, u_t, into orthogonal structural shocks, e_t. We assume that the structural shocks are zero-mean random variables with covariance matrix $E[e_t e_t'] = I_k$. The relationship between prediction errors and structural shocks can be expressed as follows:

$$u_t = F e_t, \tag{3}$$

where the $(k \times k)$ matrix F provides an exact factorization of the prediction error covariance matrix, i.e. $FF' = \Sigma_u$. For the purpose of this paper, and without loss of generality, it will be useful to decompose F into two components:

$$F = \tilde{A}Q, \tag{4}$$

where the $(k \times k)$ matrix $\tilde{A}$ is the lower-triangular Cholesky factor of Σ_u and the $(k \times k)$ matrix Q is an orthonormal square matrix satisfying $QQ' = I_k$ and $\det(Q) = 1$. The matrix Q is thus a rotation matrix. It can be constructed as the product of at most $k(k-1)/2$ Givens rotations $G_{ij}(\theta)$, each acting on a two-dimensional subspace with rotation angle θ.[4] The Givens rotation $G_{ij}(\theta)$ in

[4] Alternatively, Householder transformations have been used in the literature (see e.g. Rubio-Ramirez, Waggoner, and Zha (2010)), mainly in order to improve the efficiency of numerical algorithms.

the (i,j) coordinate plane takes the form

$$G_{ij}(\theta) = \begin{bmatrix} 1 & \dots & 0 & \dots & 0 & \dots & 0 \\ \vdots & \ddots & \vdots & & \vdots & & \vdots \\ 0 & \dots & cos(\theta) & \dots & sin(\theta) & \dots & 0 \\ \vdots & & \vdots & \ddots & \vdots & & \vdots \\ 0 & \dots & -sin(\theta) & \dots & cos(\theta) & \dots & 0 \\ \vdots & & \vdots & & \vdots & \ddots & \vdots \\ 0 & \dots & 0 & \dots & 0 & \dots & 1 \end{bmatrix}. \tag{5}$$

In this paper we use the Cholesky decomposition for a particular variable ordering because this allows for the derivation of simple analytical expressions (see Sections 3 and 4). Note, however, that the inference does not depend on the use of this particular decomposition. Any other exact factorization of Σ_u will yield the same inference (see Uhlig (2005), Appendix B). This applies both to alternative variable orderings in the Cholesky decomposition and alternative matrix decomposition methods such as the eigenvalue-eigenvector decomposition used e.g. by Canova and De Nicoló (2002) and Peersman (2005). Of course, the matrix Q will differ for alternative factor matrices $\tilde{A}$. Yet, the sign-restriction identification problem is not linked to the chosen form of $\tilde{A}$ but due to the fact that there are in general infinitely many admissible rotation matrices consistent with a given set of sign restrictions.

We impose sign restrictions only on the *impact* responses of the set (a subset) of model variables in response to a given shock. Imposing restrictions also at longer horizons reduces the appeal of analytical solutions as in this case the coefficients of the reduced-form lag polynomial $B(L)$ enter the solution. As is well-known the estimates of these coefficients tend to be rather imprecise, whereas the estimates of the error covariance matrix tend to be rather precise. Importantly, the sign of residual covariances tends to be robust to the inclusion of additional variables and/or deterministic terms, which is what we need for the derivation of our analytical results.

Our focus on contemporaneous sign restrictions is less restrictive than it may appear. Empirical studies using the sign-restriction approach typically find that the results are insensitive qualitatively to variations in the horizon at which impulse responses are restricted (see e.g. Uhlig (2005), p. 398; Peersman (2005), p. 191; Mountford and Uhlig (2009), p. 965). Moreover, as Fry and Pagan (2009) report in their review of the sign restriction approach, a consensus seems to be emerging in the literature that imposing contemporaneous restrictions only might be preferable to imposing restrictions also on longer lags.

2.3 The analytical Cholesky decomposition

The key to our analytical approach is the existence of an analytical expression for the lower-triangular Cholesky decomposition of the covariance matrix Σ_u, where we assume that the latter is real symmetric positive definite. The Cholesky

decomposition has a recursive structure, which greatly simplifies the derivation of the individual elements of the Cholesky factor matrix. The algorithms to compute the individual elements of the Cholesky factor start from the upper left corner and proceed to calculate the matrix either row by row (Cholesky-Banachiewicz algorithm) or column by column (Cholesky-Crout algorithm).

The lower-triangular Cholesky decomposition of the covariance matrix Σ_u yields the following expression for the Cholesky factor $\tilde{A}$, where to save space we report only the elements of the first and second columns:

$$\tilde{A} = \begin{bmatrix} \sigma_1 & 0 & 0 & ... & 0 \\ \sigma_2\rho_{12} & \sigma_2\sqrt{1-\rho_{12}^2} & 0 & ... & 0 \\ \vdots & \vdots & \ddots & & \vdots \\ \sigma_j\rho_{1j} & \sigma_j\frac{\rho_{2j}-\rho_{1j}\rho_{12}}{\sqrt{1-\rho_{12}^2}} & & \ddots & 0 \\ \vdots & \vdots & & \ddots & \vdots \\ \sigma_k\rho_{1k} & \sigma_k\frac{\rho_{2k}-\rho_{1k}\rho_{12}}{\sqrt{1-\rho_{12}^2}} & ... & ... & \tilde{a}_{kk} \end{bmatrix}. \tag{6}$$

This analytical expression reveals that the elements in the first column of $\tilde{A}$ obtain as $\tilde{a}_{i1} = \sigma_i\rho_{1i}$ for all i, and the elements in the second column on and below the principal diagonal obtain as $\tilde{a}_{i2} = \sigma_i\left(\rho_{2k} - \rho_{1k}\rho_{12}\right)/\sqrt{1-\rho_{12}^2}$, $i = 2, ..., k$. The elements on the principal diagonal, beyond having simpler analytical expressions than those below the main diagonal because of $\rho_{ii} = 1$ for all i, share the property that they are all strictly larger than zero ($\tilde{a}_{ii} > 0$ for all i).

3 The simplest case: bivariate VAR models

This section analyzes bivariate VAR models. This simple baseline case helps clarifying the intuition behind the sign-restriction approach. It is also sufficient for establishing a number of important results that generalize to the multivariate case as we show in Section 4.

In the following we consider a setup in which two shocks have to be identified. This setup nests typical problems macroeconomists are confronted with: think of the first shock as non-policy shock and of the second shock as policy shock. Further assume that the model consists of a non-policy variable that is ordered first in the VAR system and of a policy instrument that is ordered second. In general, the response of the non-policy variable to the policy shock will be the object of interest (F_{12} in our notation). In this setup the first shock serves to control for comovements in the two variables that are unrelated to unanticipated changes in policy but originate from non-policy shocks.

3.1 The set of pure sign-restriction solutions

To formalize the above discussion we make the following assumptions. We impose sign restrictions on three contemporaneous impulse responses and leave the

response of the first variable (the non-policy variable) to the second shock (the policy shock), F_{12}, unrestricted. Specifically, we require the impact responses of the first variable to the first shock and of the second variable to the second shock to be non-negative ($F_{11} \geq 0$, $F_{22} \geq 0$). These restrictions can be viewed as normalizations determining the sign of the shock. In addition, we require the impact response of the second variable (the policy variable) to the first shock (the non-policy shock) to be non-negative ($F_{21} \geq 0$). We view this as the case that naturally arises in many macroeconomic applications (see above). We therefore concentrate on this case. The derivations for the alternative case ($F_{21} \leq 0$) are analogous to the ones presented below.

Assumption (Sign restrictions on contemporaneous impulse responses)**.** *Let the* (2×2) *matrix* F *be an exact factor of the* (2×2) *covariance matrix* Σ_u*, i.e.* $FF' = \Sigma_u$*. We impose the following sign restrictions on the individual elements of the factor matrix:* $F_{11} \geq 0$*,* $F_{21} \geq 0$ *and* $F_{22} \geq 0$*.*

The impact impulse responses are obtained from the system of equations (4), which here is (2×2):

$$\begin{bmatrix} F_{11} & F_{12} \\ F_{21} & F_{22} \end{bmatrix} = \begin{bmatrix} \sigma_1 & 0 \\ \sigma_2\rho_{12} & \sigma_2\sqrt{1-\rho_{12}^2} \end{bmatrix} \begin{bmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{bmatrix}. \tag{7}$$

To facilitate the derivation of analytical results, without loss of generality, we express the error correlation coefficient ρ_{12} as angle.

Definition (Angle represention of the error correlation coefficient)**.** Let $\varphi \equiv \arccos(\rho_{12})$ be the angle represention of the error correlation coefficient.

Using this definition we can express $\rho_{12} = \cos(\varphi)$ and $\sqrt{1-\rho_{12}^2} = \sin(\varphi)$. Note that $\rho_{12} \in (-1, 1)$ implies $\varphi \in (0, \pi)$. The angle φ is strictly decreasing in ρ_{12}, with $\varphi \to \pi$ as $\rho_{12} \to -1$, $\varphi = \frac{\pi}{2}$ for $\rho_{12} = 0$ and $\varphi \to 0$ as $\rho_{12} \to 1$. Using the angle representation of ρ_{12}, the expressions for the impact impulse responses implied by Equation (7) simplify to (with sign restrictions shown where applicable):

$$\begin{aligned} F_{11} &= \sigma_1 \cos(\theta) \geq 0, & (8) \\ F_{21} &= \sigma_2 \cos(\theta - \varphi) \geq 0, & (9) \\ F_{12} &= -\sigma_1 \sin(\theta), & (10) \\ F_{22} &= -\sigma_2 \sin(\theta - \varphi) \geq 0, & (11) \end{aligned}$$

where we have made use of basic trigonometric identities.[5]

Definition (No-puzzle and puzzle solutions)**.** If for a pure sign-restriction solution $F_{12} \leq 0$, this solution is called a *no-puzzle solution*. If instead for a pure sign restriction solution $F_{12} \geq 0$, this solution is called a *puzzle solution*.

[5] The expression for F_{21} uses the angle difference identity $\cos(\theta)\cos(\varphi) + \sin(\theta)\sin(\varphi) = \cos(\theta - \varphi)$, while the expression for F_{22} uses the angle difference identity $\sin(\theta)\cos(\varphi) - \cos(\theta)\sin(\varphi) = \sin(\theta - \varphi)$.

Remember that the object of interest, F_{12}, is representative of e.g. the impact response of prices to a monetary policy shock or the impact response of real output to a tax shock. In such settings, $F_{12} > 0$ implies that prices increase on impact in response to a contractionary monetary policy shock ("the price puzzle") or, equivalently, that output increases on impact in response to a tax hike ("the tax multiplier puzzle"). Note that the above definition purposefully leaves open whether $F_{12} = 0$ is to be considered a puzzle or not.

We now have all the ingredients to characterize the set of all pure sign-restrictions solutions for the bivariate case.

Proposition 1 (The set of all pure sign-restriction solutions)**.** *Let S be the set of all pure sign-restriction solutions satisfying the restrictions $F_{11}(\theta) \geq 0, F_{21}(\theta) \geq 0, F_{22}(\theta) \geq 0$. Let $S_1 \subset S$ be the subset of no-puzzle solutions for which $F_{12}(\theta) \leq 0$ and let $S_2 \subset S$ be the subset of puzzle solutions for which $F_{12}(\theta) \geq 0$ such that $S = S_1 \cup S_2$.*
The subset of no-puzzle solutions, for given $\varphi \in (0, \pi)$, is

$$S_1 \equiv \left\{ \theta \in [-\pi, \pi] : \max(-\frac{\pi}{2} + \varphi, 0) \leq \theta \leq \min(\varphi, \frac{\pi}{2}) \right\}.$$

The subset of puzzle solutions, for given $\varphi \in \left(0, \frac{\pi}{2}\right]$, is

$$S_2 \equiv \left\{ \theta \in [-\pi, \pi] : -\frac{\pi}{2} + \varphi \leq \theta \leq 0 \right\}.$$

The set of all pure sign-restriction solutions, for given $\varphi \in (0, \pi)$, is

$$S \equiv \left\{ \theta \in [-\pi, \pi] : -\frac{\pi}{2} + \varphi \leq \theta \leq \min(\varphi, \frac{\pi}{2}) \right\}.$$

The set of all sign-restriction solutions, S, and the subset of no-puzzle solutions, S_1, are non-empty for all $\rho_{12} \in (-1, 1)$, i.e. for less than perfect correlation between the prediction errors.
The subset of puzzle solutions, S_2, is non-empty if and only if the correlation between the prediction errors is non-negative but less than perfect:

$$S_2 \neq \emptyset \Leftrightarrow 0 \leq \rho_{12} < 1.$$

The subsets of no-puzzle solutions, S_1, and of puzzle solutions, S_2, have equal size (and are non-empty) if and only if $\varphi = \frac{\pi}{4}$, i.e. $\rho_{12} = 1/\sqrt{2}$.
The Cholesky factor is a pure sign-restriction solution if and only if the subset of puzzle solutions is non-empty, i.e. if and only if $0 \leq \rho_{12} < 1$.

The full proof of Proposition 1 is given in Appendix A. The proof proceeds as follows: the set of all pure sign restriction solutions is a subset of the set all possible rotation angles.[6] Now, the sign restrictions put restrictions on the set

[6] We restrict attention to the range $\theta \in [-\pi, \pi]$, which covers clockwise rotations with angle between -180° and 0° and counterclockwise rotations with angle between 0° and 180°, i.e. together covers the full circle.

of admissible rotation angles. Appendix A shows that the sign restrictions imply that the set of no-puzzle solutions is equivalent to the set of counterclockwise rotations with angle between at least 0° and at most 90°, while the set of puzzle solutions is equivalent to the set of clockwise rotations with angle between at least -90° and 0°. It is also shown that the set of all pure sign restriction solutions and the subset of no-puzzle solutions are non-empty for all feasible values of ρ_{12}, i.e. solutions satisfying the sign restrictions exist under general conditions.

Figure 1 illustrates the set and subsets of pure sign restriction solutions for four alternative values of the correlation coefficient between reduced-form prediction errors ρ_{12}. The subplots show all impact responses of variable 1 (the non-policy variable) to shock 2 (the policy shock) consistent with the sign restrictions. The first row shows results for non-positive error correlation ($\rho_{12} \leq 0$). In this case the set of all pure sign restriction solutions S, represented by the curve segment AB, coincides with the subset of no-puzzle solutions S_1, represented by the curve segment CD. Also, these two sets are the smaller the more negative is the error correlation; in the limit ($\rho_{12} \to -1$) they collapse into a single element (point A). The subset of puzzle solutions S_2 is either empty (for $\rho_{12} < 0$) or consists of just one element, namely the Cholesky factor (for $\rho_{12} = 0$).

The second row of Figure 1 shows results for positive error correlation ($\rho_{12} > 0$). In this case, the interval of rotation angles characterizing the set of all pure sign restriction solutions has length $\pi/2$. Whereas the subset of no-puzzle solutions shrinks with increasing error correlation, the subset of puzzle solutions, represented by the curve segment EF, is the larger the stronger the correlation. For $\rho_{12} = 1/\sqrt{2}$ the intervals characterizing these two subsets have equal length $\pi/4$. In the limit ($\rho_{12} \to 1$) the subset of no-puzzle solutions collapses into a single element (the Cholesky factor), while the set of puzzle solutions reaches it maximum possible size, i.e. close to $\pi/2$.

In addition to the sets of pure sign restriction solutions, Figure 1 also shows two particular points that will play a prominent role in Sections 3.2 and 3.3, namely the midpoint of the set of all sign restriction solutions (point G) and the midpoint of the set of no-puzzle solutions (point H). Section 3.2 shows that the midpoint of S (point G) is equivalent to the median rotation angle typically picked by the pure sign restriction approach, while Section 3.3 shows that the midpoint of S_1 is equivalent to the "optimal" rotation angle picked by the standard penalty function approach. Figure 1 shows that points G and H differ for positive error correlation. For $\rho_{12} > 1/\sqrt{2}$ the midpoint of S is an element of the set of puzzle solutions.

Proposition 1 establishes a simple condition for existence of puzzle solutions based on the sign of the correlation coefficient between the two prediction errors. Puzzle solutions exist if and only if this correlation is non-negative (and less than perfect). And, interestingly, only when this condition is fulfilled, is the Cholesky factor a pure sign restriction solution.

The implications for the applied macroeconometric literature are straightforward: the "price puzzle" is consistent with a given reduced-form VAR model whenever the correlation between price residuals and interest rate residuals is non-negative. This also explains why attempts to resolve the price puzzle some-

times have been successful but sometimes not. Sims (1992) noted that adding commodity prices to a VAR model has the potential to resolve the price puzzle, whereas subsequent studies for US data show that this is not in general true but holds only for specific subsamples (see e.g. Hanson (2004)). Our analysis suggests a simple explanation for these findings: the inclusion of commodity prices resolves the price puzzle if and only if it "sufficiently" reduces the correlation between price and interest rate residuals.[7] The reduction may be different for different subsamples, explaining why the price puzzle is resolved for some subsamples but not for others.

The subset of "tax multiplier puzzle" solutions, on the other hand, is very unlikely to be empty at all for typical empirical applications. Regardless of the specification of the reduced-form VAR model, prediction errors for taxes and for output are likely to exhibit strong positive correlation, if – as is typically the case – output is a good proxy for the macroeconomic tax base. As a consequence, both the subset of no-puzzle and the subset of puzzle solutions will be non-empty. A natural question then is whether the solution(s) an applied econometrician picks will be from the subset of no-puzzle solutions or the subset of puzzle solutions or both. In the applied literature two approaches have been used to report results for the sign restriction approach: first, the pure sign restriction approach, and, second, the penalty function approach. The next two subsections investigate the properties of these two approaches in turn.

3.2 The analytics of the pure sign restriction approach

Given that the set of pure sign restriction solutions in most cases contains a very large number of elements with potentially widely differing implications, the question arises how to summarize the information contained in this set.[8] The empirical literature applying the pure sign restriction approach usually applies a numerical algorithm similar to the following (see e.g. Canova (2007), p. 151):

1. Conditional on the estimate of a reduced-form model draw rotation angles from a uniform distribution.

[7] We say "sufficiently" because in empirical applications it may be enough for the correlation to be low but still positive for the price puzzle to disappear. The reason is that the subset of no-puzzle solutions is non-empty in any case. Whether an applied researcher picks a solution from the subset of puzzle solutions or from the subset of no-puzzle solutions depends not only on the relative size of these two subsets but also on the criterion used to pick one particular solution from the set of all pure sign restriction solutions. The criteria used in the literature are investigated in Sections 3.2 and 3.3. Of course, if the subset of puzzle solutions is empty, the results reported will trivially come from the subset of no-puzzle solutions.

[8] It has to be stressed that we deal exclusively with structural model uncertainty and disregard the sampling uncertainty surrounding the reduced-form VAR model (see Fry and Pagan (2009) for an in-depth discussion). For the derivation of the analytical results we treat the estimate of the reduced-form model as fixed. For our purposes it does not matter whether this estimate is an OLS estimate (frequentist approach) or a draw from the posterior distribution (Bayesian approach). What we are after is the set of structural models consistent with a given estimate of the reduced-form model and given sign restrictions.

2. Compute the impulse responses for each draw and check whether they satisfy the sign restrictions. If they do, accept the draw, else reject the draw.

3. Repeat steps 1 and 2 until the desired number of draws satisfying the restrictions are found.

4. Report summary statistics.

The literature differs with regard to the summary statistics reported. Some studies report the minimum and the maximum of the impulse responses for the set of draws as an estimate for the bounds of the interval of impulse responses satisfying the sign restrictions (see e.g. Uhlig (2005)), while other studies report the median point of the interval (see e.g. Peersman (2005)).

These differences in procedure can have profound implications for the results. Whereas the first procedure is likely to yield a range that encompasses both no-puzzle solutions and puzzle solutions (if the latter set is non-empty), the second procedure singles out one element of either the set of no-puzzle solutions or the set of puzzle solutions.

Our analytical approach allows us to make theoretical predictions on any summary statistics that may be of interest, without the need for resorting to simulation methods. The bounds of the interval of impulse responses satisfying the sign restrictions, e.g., are simply given by the bounds of the set of all pure sign restriction solutions S as shown in Proposition 1. We summarize the most interesting results on the midpoint of this interval, which for the uniform distribution is equivalent to its median and mean, in Proposition 2.

Proposition 2 (Properties of the pure sign restriction approach)**.** *Let $\theta|\,\Sigma_u$ be uniformly distributed with support $\theta \in S$, i.e. the set of all pure sign-restriction solutions (see Proposition 1).*
Then, the median rotation angle $\bar{\theta}$ is equal to the midpoint of S:

$$\bar{\theta} = \begin{cases} -\frac{\pi}{4} + \varphi & for\ given\ \varphi \in (0, \frac{\pi}{2}], \\ \frac{\varphi}{2} & for\ given\ \varphi \in [\frac{\pi}{2}, \pi). \end{cases}$$

For $\varphi \in [\frac{\pi}{2}, \pi)$, i.e. $\rho_{12} \in (-1, 0]$, $\bar{\theta}$ is equal to the midpoint of the subset of no-puzzle solutions S_1 (Point H in Figure 1). For $\varphi \in (0, \frac{\pi}{2}]$, i.e. $\rho_{12} \in [0, 1)$, $\bar{\theta}$ is equal to the sum of the midpoints of the subset of no-puzzles solutions S_1 and the subset of puzzle solutions S_2 (Point G in Figure 1).
For $\varphi \in (\frac{\pi}{4}, \pi)$, i.e. *$\rho_{12} \in (-1, \frac{1}{\sqrt{2}})$, $\bar{\theta}$ is an element of the subset of no-puzzle solutions S_1. For $\varphi \in (0, \frac{\pi}{4})$,* i.e. *$\rho_{12} \in (\frac{1}{\sqrt{2}}, 1)$, $\bar{\theta}$ is an element of the subset of puzzle solutions S_2. For $\varphi = \frac{\pi}{4}$,* i.e. *$\rho_{12} = \frac{1}{\sqrt{2}}$, the median rotation angle $\bar{\theta}$ is equal to 0, i.e. the pure sign restriction approach picks the Cholesky factor.*

Proof. The median, mean and midpoint of a uniformly distributed random variable x with support $x \in [a, b]$ are all equal to the average of the minimum and maximum of its support, i.e. $\bar{x} = \frac{1}{2}(a+b)$. Applying this property of the uniform

distribution to the set of pure sign restriction solutions S and its subsets S_1 and S_2 yields the expressions for the midpoints given in Proposition 2. □

Proposition 2 highlights one result that we believe is important for applied macroeconomic research: an econometrician relying merely on the median of the distribution of pure sign restriction solutions will not report puzzling results unless the coefficient of correlation between reduced-form prediction errors exceeds $1/\sqrt{2}$ (around 0.71). Yet, we know from Proposition 1 that the set of puzzle solutions can be quite large also for smaller values of the correlation coefficient. It, thus, seems advisable to report the bounds of the interval of impulse responses satisfying the sign restrictions in addition to the median in order to get the full picture of structural models consistent with a given reduced-form model.[9] As the next section shows the same reasoning applies to the penalty function approach.

3.3 The analytics of the penalty function approach

As an alternative to the pure sign restriction approach, the literature has used the penalty function approach to select one particular solution out of the set of pure sign restriction solutions (see e.g. Faust (1998), Uhlig (2005) and Mountford and Uhlig (2009)). More generally, the penalty function approach is used to numerically solve constrained nonlinear optimization problems for which closed-form analytical solutions may be hard to obtain or may not exist at all (see Judd (1998), pp. 123-25). The idea is to replace the constraints – here: the sign restrictions – with a continuous penalty function that permits but (heavily) penalizes choices that violate the constraints. As a result, the constrained problem is replaced with an unconstrained one.

In this subsection we show that it is possible to analytically solve the nonlinear optimization problem underlying the sign-restriction penalty function approach for the bivariate case. We propose a general objective function that nests the standard penalty function used in the literature. Compared to the standard penalty function, this objective function introduces a single additional parameter, whose value determines which pure sign restriction solution is picked as the optimum. Moreover, we show that depending on which impulse responses enter as arguments in the objective function the optimum will be either in the set of no-puzzle solutions or in the set of puzzle solutions. We show that the standard penalty function used in the literature implies that the optimum is in the set of no-puzzle solutions. We formally prove what should be intuitively clear: the standard penalty function is an identifying assumption.

The standard penalty function has the sum of some or all impact impulse responses to a given shock as its arguments, where in general the impulse response

[9] In the literature sometimes "confidence intervals" are reported for rotation angles. We agree with Paustian (2007) and Fry and Pagan (2009) that there is little justification for this procedure because *all* rotation angles consistent with a given set of sign restrictions are equally likely. The likelihood function does not contain information on the translation from reduced to structural form. This is the essence of the identification problem.

of variable i to shock j is scaled by the standard deviation of this variable's one-step-ahead prediction error, σ_i.[10] In our notation, the individual arguments of the penalty function can thus be written as F_{ij}/σ_i. This expression is equivalent to the square root of the fraction of variable i´s one-step-ahead forecast error variance explained by shock j.

The literature, in general, proceeds as follows: first, numerically minimize a penalty function having some or all impulse responses to the first shock as its arguments. Second, if a second shock has to be identified, minimize a penalty function having some impulse responses to the second shock as arguments, imposing the further constraint that the second shock is orthogonal to the first shock. The identification of further shocks proceeds analogously. In general, all shocks are required to be orthogonal to the shocks previously identified.

In the bivariate case, the problem simplifies as there is a maximum of two shocks. There can be only one penalty function as the second shock imperatively has to explain all fluctuation not explained by the first shock. In line with the literature, we here assume that the penalty function has the impulse responses of the first and second variables to the first shock as its arguments. Then, the following general objective function nests the standard penalty function:

$$\Omega_1 \equiv \frac{F_{11}}{\sigma_1} + b\frac{F_{21}}{\sigma_2} = \cos(\theta) + b\cos(\theta - \varphi), \tag{12}$$

where the parameter $b \in \mathbb{R}_+$ gives the weight assigned to the impact impulse response of the second variable to the first shock. The standard penalty function used in the literature obtains for $b = 1$, i.e. for equal weights on the scaled impulse responses. Two further special cases are of particular interest: For $b = 0$ the objective function reduces to $\Omega_1 = F_{11}/\sigma_1$, i.e. the optimization problem reduces to finding the rotation angle θ for which the fraction of the first variable's one-step-ahead forecast error variance explained by the first shock is maximized. For $b \to \infty$ the objective function reduces to $\Omega_1 \to F_{21}/\sigma_2$, i.e. the optimization problem reduces to finding the rotation angle θ for which the fraction of the second variable's one-step-ahead forecast error variance explained by the first shock is maximized.

Three characteristics of this optimization problem greatly facilitate its analytical solution. First, expressing the problem in terms of trigonometric functions implies that the orthonormality restriction on the rotation matrix Q is automatically satisfied because $\cos^2(\omega) + \sin^2(\omega) = 1$ holds for any $\omega \in \mathbb{R}$. Second, since we can analytically characterize the set of all pure-sign restriction solutions, S, (see Proposition 1), we know that the domain of Ω_1 is a closed and bounded interval on $\mathbb{R}$. Third, since Ω_1 is a continuous and concave function on S,[11] we have by Weierstrass's theorem that Ω_1 achieves its global maximum on

[10] Recall that we restrict our attention to contemporaneous sign restrictions.

[11] The cosine function is continuous for all real numbers. To show that Ω_1 is concave on S it is sufficient to show that $\cos(\theta)$ and $\cos(\theta - \varphi)$ in Equation (12) are both concave on S since b is a nonnegative number (see Simon and Blume (1994), Theorem 21.8, p. 519). Using the second-derivative test it is straightforward to show that $\cos(\theta)$ is concave on the interval $[-\frac{\pi}{2}, \frac{\pi}{2}]$ and $\cos(\theta - \varphi)$ is concave on the interval $[-\frac{\pi}{2} + \varphi, \frac{\pi}{2} + \varphi]$. The set of all pure sign-restriction solutions S is a subset of both intervals for all $\varphi \in (0, \pi)$.

its domain. Furthermore, this global maximum is unique.[12] In sum, we neither have to explicitly account for the equality constraint implied by the orthonormality assumption, nor for the inequality constraints associated with the sign restrictions. Both sets of constraints will be automatically satisfied.

We now are in a position to characterize the analytical solution to the optimization problem underlying the penalty function approach.

Proposition 3 (The optimization problem of the penalty function approach)**.** *Let $b \in \mathbb{R}_+$ be a parameter. Let $\Omega_1 \equiv \frac{F_{11}}{\sigma_1} + b\frac{F_{21}}{\sigma_2}$ be the concave objective function to be maximized, defined on the convex subset $\theta \in [-\frac{\pi}{2} + \varphi, \min(\varphi, \frac{\pi}{2})]$ of $\mathbb{R}$ given by the set of all pure sign-restriction solutions S (see Proposition 1).*

Then, for given b and $\varphi \in (0, \pi)$, $\theta^*(b, \varphi)$ is the global maximizer of Ω_1 on S. The global maximizer is an element of the subset of no-puzzle solutions, $\theta^*(b, \varphi) \in S_1$, for all $b \in \mathbb{R}_+$ and all $\varphi \in (0, \pi)$.

Case 1: Non-negative error correlation ($\rho_{12} \geq 0$), i.e. $\varphi \in (0, \frac{\pi}{2}]$.
The global maximum of Ω_1 on S, for given $\varphi \in (0, \frac{\pi}{2}]$, obtains for

$$\theta^*(b, \varphi) = \arctan\left(\frac{\sin\varphi}{\frac{1}{b} + \cos\varphi}\right).$$

The optimal angle of rotation $\theta^*(b, \varphi) \in [0, \varphi]$ and is monotonically increasing in b. For $b \in (0, \infty)$ the maximum is an interior maximum, for $b = 0$ and $b \to \infty$ the maximum occurs on the boundary of S_1.
Special cases:
For $b = 0$, the maximum obtains for $\theta^* = 0$, i.e. the optimum coincides with the Cholesky factor, i.e. the lower bound of S_1 (Point B in Figure 1).
For $b \to \infty$, the maximum obtains for $\theta^* = \varphi$, i.e. the optimum coincides with the upper bound of S_1 (Point A in Figure 1).
For $b = 1$ (the standard case in the literature), the maximum obtains for $\theta^* = \frac{\varphi}{2}$, *i.e. the optimum coincides with the midpoint of S_1 (Point H in Figure 1).*

Case 2: Negative error correlation ($\rho_{12} < 0$), i.e. $\varphi \in (\frac{\pi}{2}, \pi)$
The global maximum of Ω_1 on S, for given $\varphi \in (\frac{\pi}{2}, \pi)$, obtains for

$$\theta^*(b, \varphi) = arccot\left(\frac{\frac{1}{b} + \cos\varphi}{\sin\varphi}\right) \iff \underline{b} < b < \overline{b},$$

where $0 < \underline{b} = -\cos(\varphi) = -\rho_{12} < 1$ and $1 < \overline{b} = -\frac{1}{\cos(\varphi)} = -\frac{1}{\rho_{12}} < \infty$.
For $b \leq \underline{b}$ and $b \geq \overline{b}$ the maximum occurs on the boundary of S_1.

[12] To establish uniqueness we need to show that Ω_1 is strictly concave $\left(\frac{d^2\Omega_1}{d\theta^2} < 0\right)$ on its entire domain. Note, first, that $\cos(\theta)$ and $\cos(\theta - \varphi)$ have continuous derivatives with respect to θ of every order, implying that Ω_1 is also infinitely continuously differentiable. Concavity implies that $\frac{d^2\cos(\theta)}{d\theta^2} = -\cos(\theta) \leq 0$ on S and $\frac{d^2\cos(\theta-\varphi)}{d\theta^2} = -\cos(\theta - \varphi) \leq 0$ on S for all $\varphi \in (0, \pi)$. Furthermore, $\frac{d^2\cos(\tilde{\theta})}{d\theta^2} = \frac{d^2\cos(\tilde{\theta}-\varphi)}{d\theta^2} = 0$ for identical angle of rotation $\tilde{\theta} \in S$ if and only if $\varphi = \pi$ (perfect negative correlation), which is ruled out by assumption. Thus, Ω_1 is strictly concave on S for all $\varphi \in (0, \pi)$, ensuring uniqueness of its global maximum.

For $b \leq \underline{b}$ the maximum obtains for $\theta^* = \varphi - \frac{\pi}{2}$, i.e. the optimum coincides with the lower bound of S_1 (Point B in Figure 1).
For $b \geq \bar{b}$ the maximum obtains for $\theta^* = \frac{\pi}{2}$, i.e. the optimum coincides with the upper bound of S_1 (Point A in Figure 1).
Special case:
For $b = 1$ (the standard case in the literature), the maximum obtains for $\theta^* = \frac{\varphi}{2}$, *i.e. the optimum coincides with the midpoint of S_1 (Point H in Figure 1).*

The proof of Proposition 3 is given in Appendix B.

Proposition 4 (Alternative interpretation of the standard penalty function)**.** *The standard penalty function selects the pure sign-restriction solution that maximizes the fraction of one-step-ahead forecast error covariance explained by the first shock if the error correlation is positive and minimizes it if the error correlation is negative (Point H in Figure 1).*

Let $\tilde{\Omega}_1 \equiv \frac{F_{11}F_{21}}{\sigma_{12}} = \frac{\cos(\theta)\cos(\theta-\varphi)}{\cos(\varphi)}$ be the objective function to be optimized, defined on the convex subset $\theta \in [\max(-\frac{\pi}{2}+\varphi, 0), \min(\varphi, \frac{\pi}{2})]$ of $\mathbb{R}$ given by the set of no-puzzle sign-restriction solutions S_1. Then, for given $\varphi \in (0,\pi) \setminus \{\frac{\pi}{2}\}$, the critical point of $\tilde{\Omega}_1$ on S_1 coincides with the global maximum of the standard penalty function $\Omega_1(b=1)$ on S (see Proposition 3):

$$\theta^*(\varphi) = \frac{\varphi}{2}.$$

For given $\varphi \in (0, \frac{\pi}{2})$, i.e. positive error correlation, the objective function $\tilde{\Omega}_1$ is strictly concave and the critical point is the unique global maximum of $\tilde{\Omega}_1$ on S_1; for given $\varphi \in (\frac{\pi}{2}, \pi)$, i.e. negative error correlation, the objective function $\tilde{\Omega}_1$ is strictly convex and the critical point is the unique global minimum of $\tilde{\Omega}_1$ on S_1.

The proof of Proposition 4 is given in Appendix C.

Proposition 5 (Pure versus standard penalty function approach)**.** *Let $\bar{\theta} = \min(-\frac{\pi}{4}+\varphi, \frac{\varphi}{2})$ be the summary solution (the median rotation angle) picked by the pure sign restriction approach (see Proposition 2). Let $\theta^* = \frac{\varphi}{2}$ be the optimal solution picked by the penalty function approach with standard parametrization $b = 1$ (see Proposition 3).*
Then, for positive correlation between the prediction errors ($\varphi \in (0, \frac{\pi}{2})$), the rotation angle picked by the penalty function approach is strictly larger than the rotation angle picked by the pure sign restriction approach:

$$\theta^* > \bar{\theta} \iff 0 < \rho_{12} < 1.$$

For non-positive error correlation ($\varphi \in [\frac{\pi}{2}, \pi)$) the two rotation angles are identical:

$$\theta^* = \bar{\theta} \iff -1 < \rho_{12} \leq 0.$$

4 The multivariate case: seemingly unrestricted variables

TO BE ADDED

5 Empirical application: the fiscal multiplier debate

TO BE ADDED

6 Conclusion

TO BE ADDED

References

Blanchard, O., and R. Perotti (2002): "An Empirical Characterization of the Dynamic Effects of Changes in Government Spending and Taxes on Output," *The Quarterly Journal of Economics*, 117(4), 1329–1368.

Canova, F. (2007): *Methods for Applied Macroeconomic Research*. Princeton, NJ: Princeton University Press.

Canova, F., and G. De Nicoló (2002): "Monetary Disturbances Matter for Business Fluctuations in the G-7," *Journal of Monetary Economics*, 49(6), 1131–1159.

Eichenbaum, M., and C. L. Evans (1995): "Some Empirical Evidence on the Effects of Shocks to Monetary Policy on Exchange Rates," *Quarterly Journal of Economics*, 110(4), 975–1009.

Faust, J. (1998): "The Robustness of Identified VAR Conclusions about Money," *Carnegie-Rochester Conference Series on Public Policy*, 49(Dec), 207–244.

Fry, R., and A. Pagan (2009): "Sign Restrictions in Structural Vector Autoregressions: A Critical Review," Unpublished manuscript.

Galí, J., J. D. López-Salido, and J. Vallés (2007): "Understanding the Effects of Government Spending on Consumption," *Journal of the European Economic Association*, 5(1), 227–270.

Golub, G. H., and C. F. van Loan (1996): *Matrix Computations*. Baltimore: Johns Hopkins University Press, 3rd edn.

Grilli, V., and N. Roubini (1996): "Liquidity Models in Open Economies: Theory and Empirical Evidence," *European Economic Review*, 40(3-5), 847–859.

Hanson, M. S. (2004): "The "Price Puzzle" Reconsidered," *Journal of Monetary Economics*, 51(7), 1385–1413.

Judd, K. L. (1998): *Numerical Methods in Economics*. Cambridge, MA: MIT Press.

Leeper, E. M., and D. B. Gordon (1992): "In Search of the Liquidity Effect," *Journal of Monetary Economics*, 29(3), 341–369.

Mountford, A., and H. Uhlig (2009): "What are the Effects of Fiscal Policy Shocks?," *Journal of Applied Econometrics*, 24(6), 960–992.

Paustian, M. (2007): "Assessing Sign Restrictions," *The B.E. Journal of Macroeconomics (Topics)*, 7(1), Article 23.

Peersman, G. (2005): "What Caused the Early Millennium Slowdown? Evidence based on Vector Autoregressions," *Journal of Applied Econometrics*, 20(2), 185–207.

Perotti, R. (2005): "Estimating the Effects of Fiscal Policy in OECD Countries," CEPR Discussion Paper 4842, Centre for Economic Policy Research.

Romer, C. D., and D. H. Romer (2010): "The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks," *American Economic Review*, forthcoming.

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*, forthcoming.

Simon, C. P., and L. Blume (1994): *Mathematics for Economists*. New York: W.W. Norton & Company.

Sims, C. (1980): "Macroeconomics and Reality," *Econometrica*, 48(1), 1–48.

Sims, C. A. (1992): "Interpreting the Macroeconomic Time Series Facts : The Effects of Monetary Policy," *European Economic Review*, 36(5), 975–1000.

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

Appendix

A Proof of Proposition 1

Proof. Abstracting from the sign restrictions, note that the interval $[-\pi, \pi]$ describes the set of all possible rotation angles, with angles measured in radians. From equation (10) we have that $F_{12} \leq 0$ if and only if $\theta \in [0, \pi]$, i.e. S_1 must be a subset of the set of counterclockwise rotations with angle between 0° and 180°. We also have that $F_{12} \geq 0$ if and only if $\theta \in [-\pi, 0]$, i.e. S_2 must be a subset of the set of clockwise rotations with angle between 0° and -180°.
The sign restrictions further restrict the set of admissible rotation angles:

$$F_{11} \geq 0 \quad \Leftrightarrow \quad -\frac{\pi}{2} \leq \theta \leq \frac{\pi}{2}, \tag{13}$$

$$F_{21} \geq 0 \quad \Leftrightarrow \quad -\frac{\pi}{2} + \varphi \leq \theta \leq \frac{\pi}{2} + \varphi, \tag{14}$$

$$F_{22} \geq 0 \quad \Leftrightarrow \quad -\pi + \varphi \leq \theta \leq \varphi. \tag{15}$$

The sign restriction on F_{11} restricts the set of admissible rotations to counterclockwise rotations with angle between 0° and 90° ($0 \leq \theta \leq \pi/2$) and clockwise rotations with angle between 0° and -90° ($-\pi/2 \leq \theta \leq 0$), i.e. sign restriction solutions fall into the first and fourth quadrants of the Cartesian plane.

Consider first the further implications of the sign restrictions for S_1. For $\varphi \leq \frac{\pi}{2}$, i.e. $\rho_{12} \geq 0$, the sign restrictions on F_{21} and F_{22} do not affect the lower bound of S_1, which is 0 in this case. However, the sign restriction on F_{22} is a binding restriction on the upper bound of S_1, which is equal to φ in this case, i.e. admissible rotation angles are strictly smaller than 90° for $\varphi < \frac{\pi}{2}$. For $\varphi > \frac{\pi}{2}$, i.e. $\rho_{12} < 0$, the sign restrictions on F_{21} and F_{22} do not affect the upper bound of S_1, which is equal to $\frac{\pi}{2}$ in this case. However, the sign restriction on F_{21} is a binding restriction on the lower bound of S_1, which is $-\frac{\pi}{2} + \varphi > 0$ in this case. Thus, the sign restrictions together imply that the set of no-puzzle solutions S_1, i.e. the set of admissible counterclockwise rotations, is described by $\theta \in [\max(-\frac{\pi}{2} + \varphi, 0), \min(\varphi, \frac{\pi}{2})]$ for all $\varphi \in (0, \pi)$. This set is non-empty for all $\varphi \in (0, \pi)$ as $\min(\varphi, \frac{\pi}{2}) \geq \max(-\frac{\pi}{2} + \varphi, 0)$. It is empty if and only if $\varphi = 0$, i.e. $\rho_{12} = 1$, or $\varphi = \pi$, i.e. $\rho_{12} = -1$, but these two cases are ruled out by assumption as otherwise Σ_u is not positive definite and the Cholesky factor does not exist.

Consider next the further implications of the sign restrictions for S_2. For $\varphi > \frac{\pi}{2}$, i.e. $\rho_{12} < 0$, the sign restriction on F_{21} implies that $0 < \theta \leq \frac{\pi}{2}$, which is inconsistent with the requirement that the set of puzzle solutions must be a subset of the set of clockwise rotations with angle between 0° and -180°; i.e. in this case the set of puzzle solutions is empty. For $\varphi = \frac{\pi}{2}$, i.e. $\rho_{12} = 0$, the set of puzzle solutions contains a single element: the Cholesky factor, which obtains for an angle of rotation equal to 0° ($\theta = 0$). For $\varphi < \frac{\pi}{2}$, i.e. $\rho_{12} > 0$, the sign restriction on F_{21} is the binding restriction on the lower bound of S_2, which is $-\frac{\pi}{2} + \varphi \leq 0$ in this case, whereas the upper bound of S_2 is 0 in this case. Thus,

the sign restrictions together imply that the set of puzzle solutions S_2, i.e. the set of admissible clockwise rotations, is described by $\theta \in [-\frac{\pi}{2}+\varphi, 0]$ and defined for $\varphi \in (0, \frac{\pi}{2}]$, i.e. $\rho_{12} \in [0, 1)$.

The set of all pure sign-restriction solutions is the union of the two subsets, i.e. $S = S_1 \cup S_2$, which gives the expression for S stated in Proposition 1.

When do S_1 and S_2 have equal size (and are non-empty)? The impulse response $F_{21}(\theta)$ is an ellipse centred at the origin and with major axis along the x-axis. Given that the ellipse is symmetric about its center, the size of the sets of pure sign restriction solutions is proportional to the length of the interval of admissible rotation angles θ. For $\varphi \in (0, \frac{\pi}{2}]$ the interval defining S_1 has length φ, and the interval defining S_2 has length $\frac{\pi}{2} - \varphi$. Equating these expressions reveals that S_1 and S_2 have equal size for $\varphi = \frac{\pi}{4}$, i.e. $\rho_{12} = 1/\sqrt{2}$. For $\varphi < \frac{\pi}{4}$, i.e. $\rho_{12} > 1/\sqrt{2}$, S_2 is larger than S_1, i.e there are more puzzle solutions than no-puzzle solutions. Instead, for $\varphi > \frac{\pi}{4}$, i.e. $\rho_{12} < 1/\sqrt{2}$, subset S_1 is larger than subset S_2.

Finally, to assess whether the Cholesky factor is a pure sign-restriction solution only the sign of $\tilde{a}_{21} = \sigma_2 \rho_{12}$ needs to be checked because the elements on its principal diagonal are positive (see Section 2.1), thereby satisfying the non-negativity restrictions on F_{11} and F_{22}. With $\sigma_2 > 0$ by definition, the Cholesky factor satisfies the sign restriction on F_{21} if and only if ρ_{12} satisfies the sign restriction. This is the case for $\rho_{12} \in [0, 1)$, which is also the condition ensuring that the set of puzzle solutions is non-empty. □

B Proof of Proposition 3

Proof. **TO BE ADDED** □

C Proof of Proposition 4

Proof. The objective function $\tilde{\Omega}_1$ is not defined for $\varphi = \frac{\pi}{2}$, i.e. for zero error correlation, as in this case the denominator of $\tilde{\Omega}_1$ is equal to 0. For given $\varphi \in (0, \pi) \setminus \{\frac{\pi}{2}\}$, the first derivative of $\tilde{\Omega}_1$ with respect to θ is

$$\frac{d\tilde{\Omega}_1}{d\theta} = -\frac{\sin(2\theta - \varphi)}{\cos(\varphi)}.$$

This expression is equal to 0 on S_1 for $\theta^* = \frac{\varphi}{2}$, which is the unique critical point of $\tilde{\Omega}_1$ on S_1. This critical point coincides with the global maximum of Ω_1 for $b = 1$ (see Proposition 3). The further properties of this critical point depend on the sign of the error correlation coefficient ρ_{12}. Note first that the objective function $\tilde{\Omega}_1$ is (more than) twice continuously differentiable. Its second derivative with respect to θ is given by

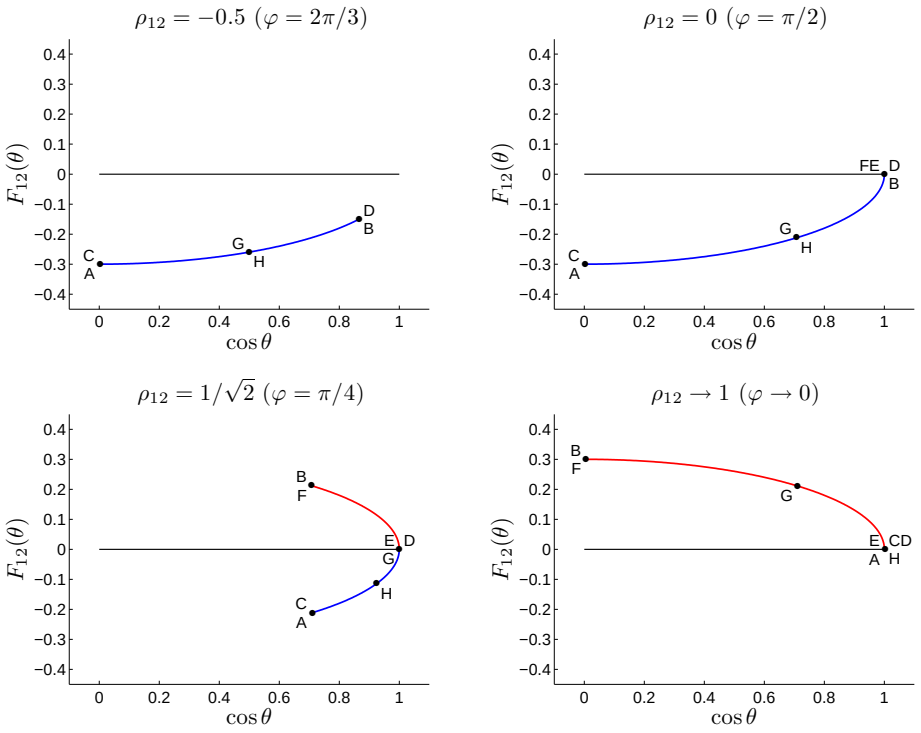
$$\frac{d^2\tilde{\Omega}_1}{d\theta^2} = -2\frac{\cos(2\theta - \varphi)}{\cos(\varphi)}.$$

The critical point of $\tilde{\Omega}_1$ is its unique global maximum on S_1 if $\tilde{\Omega}_1$ is strictly concave $\left(\frac{d^2\tilde{\Omega}_1}{d\theta^2} < 0\right)$ on S_1, whereas it is its unique global minimum if $\tilde{\Omega}_1$ is strictly convex$\left(\frac{d^2\tilde{\Omega}_1}{d\theta^2} < 0\right)$ on S_1 (Simon and Blume (1994), Theorem 3.6, p. 56).

Consider first the case of positive error correlation, i.e. $\varphi \in (0, \frac{\pi}{2})$. The second derivative of $\tilde{\Omega}_1$ is strictly negative if and only if $\cos(2\theta - \varphi) > 0$ as $\cos(\varphi) > 0$ for all $\varphi \in (0, \frac{\pi}{2})$. The expression $\cos(2\theta - \varphi) > 0$ for all $\theta \in \left(\frac{1}{2}(-\frac{\pi}{2} + \varphi), \frac{1}{2}(\frac{\pi}{2} + \varphi)\right)$. The set of no-puzzle solutions S_1, for which $\theta \in [0, \varphi]$, is a subset of this interval for all $\varphi \in (0, \frac{\pi}{2})$, proving that $\tilde{\Omega}_1$ is strictly concave on S_1 for all $\varphi \in (0, \frac{\pi}{2})$.

Consider next the case of negative error correlation, i.e. $\varphi \in (\frac{\pi}{2}, \pi)$. The second derivative of $\tilde{\Omega}_1$ is strictly positive if and only if $\cos(2\theta - \varphi) > 0$ as $\cos(\varphi) < 0$ for all $\varphi \in (\frac{\pi}{2}, \pi)$. Again, the expression $\cos(2\theta - \varphi) > 0$ for all $\theta \in \left(\frac{1}{2}(-\frac{\pi}{2} + \varphi), \frac{1}{2}(\frac{\pi}{2} + \varphi)\right)$, and the set of no-puzzle solutions S_1, for which $\theta \in [-\frac{\pi}{2} + \varphi, \frac{\pi}{2}]$, is a subset of this interval for all $\varphi \in (\frac{\pi}{2}, \pi)$, proving that $\tilde{\Omega}_1$ is strictly convex on S_1 for all $\varphi \in (\frac{\pi}{2}, \pi)$. □

Figure 1: The set of pure sign restriction solutions: Impact response of variable 1 to shock 2 ($F_{12}(\theta)$) for alternative ρ_{12}



Note: Each subplot shows the impact response of variable 1 to shock 2 for a different value of ρ_{12} (φ) for all rotation angles θ satisfying the sign restrictions. The plots build on the assumption $\sigma_1 = 0.3$ (σ_1 determines the vertical compression/stretch of the ellipse relative to the unit circle). The curve segment AB represents the set of all pure sign restriction solutions S (with upper bound A and lower bound B); the segment CD represents the subset of no-puzzle solutions S_1 (upper bound C, lower bound D); the segment EF represents the subset of puzzle solutions S_2 (upper bound E, lower bound F). The point G represents the midpoint of S (the median response for the pure sign restriction approach; see Section 3.2); the point H represents the midpoint of S_1 (the optimum of the standard penalty function; see Section 3.3).