Federal Reserve Bank of New York
Staff Reports

Identifying Shocks via Time-Varying Volatility

Daniel J. Lewis

Staff Report No. 871
October 2018
Revised May 2019

Identifying Shocks via Time-Varying Volatility
Daniel J. Lewis
Federal Reserve Bank of New York Staff Reports, no. 871
October 2018; revised May 2019
JEL classification: C32, C58, E20, E62, H30

Abstract

An *n*-variable structural vector auto-regression (SVAR) can be identified (up to shock order) from the evolution of the residual covariance across time if the structural shocks exhibit heteroskedasticity (Rigobon (2003), Sentana and Fiorentini (2001)). However, the path of residual covariances can only be recovered from the data under specific parametric assumptions on the variance process. I propose a new identification argument that identifies the SVAR up to shock orderings using the autocovariance structure of second moments of the residuals, implied by an arbitrary stochastic process for the shock variances. These higher moments are available without parametric assumptions like those required by existing approaches. The conditions required for identification can be tested using a simple procedure. The identification scheme performs well in simulations. I apply the approach to the debate on fiscal multipliers and obtain estimates lower than those of Blanchard and Perotti (2002) and Mertens and Ravn (2014), but in line with more recent studies.

Key words: identification, impulse response function, structural shocks, SVAR, fiscal multiplier, time-varying volatility, heteroskedasticity

---

Lewis: Federal Reserve Bank of New York (email: daniel.lewis@ny.frb.org). The author thanks Jim Stock, Gabriel Chodorow-Reich, Neil Shephard, Mikkel Plagborg-Møller, Lutz Kilian, Jón Steinsson, Giri Parameswaran, Domenico Giannone, Richard Crump, Andrew Patton, and participants at various seminars for helpful comments.

1 Introduction

The central challenge of structural vector autoregression (SVAR) analysis is to identify underlying structural shocks from observable VAR innovations (one-step ahead reduced-form forecast errors). For example, an innovation to tax revenues could represent either a true tax shock or the effect of automatic stabilizers as a response to changing macroeconomic conditions. Policy analysis centers on the relationships between structural shocks and observables. In an SVAR, the reduced-form innovations, η_t, are expressed as a linear combination of the underlying shocks, ε_t: $\eta_t = H\varepsilon_t$ for some contemporaneous response matrix H. Up to second moments, these equations have a multiplicity of solutions for H; economic assumptions are generally needed for identification. The majority of approaches use "internal instruments", restricting elements of H to identify the remainder. These restrictions can be short-run exclusions (Sims (1980)), long-run exclusions (Blanchard & Quah (1986)), on signs (Uhlig (2005)), or calibrated parameters (Blanchard & Perotti (2002)). More recently, "external instruments" have been proposed as an alternative, as in Mertens & Ravn (2013). However, many of these assumptions are frequently controversial.

A smaller literature offers identification based on statistical properties of the innovations. Sentana & Fiorentini (2001) and Rigobon (2003) share the important insight that if the variances of the structural shocks change over time, shocks can be identified from the reduced-form covariances at different points in time. However, this path of reduced form covariances can be recovered by the econometrician only under specific parametric models. Rigobon's (2003) method fits discrete variance regimes to the data, either based on external information or estimation. Sentana & Fiorentini (2001) use the full path of covariances, recoverable from the data only under models like Generalized Autoregressive Conditional Heteroskedasticity (GARCH).[1] Generalizations have been made to Markov switching (Lanne, Lütkepohl, & Maciejowska (2010)) and smooth transitions between regimes (Lütkepohl & Netšunajev (2017)). All of these approaches rely on knowledge of the path of variances over time and thus parametric features allowing that path to be consistently estimated, which has so far limited researchers to choose one of the few models that can be accommodated. There is compelling evidence of time-varying volatility in US macroeconomic aggregates, as documented by Stock & Watson (2002), Blanchard & Simon (2001), and Jurado, Ludvigson, & Ng (2015), so identification based on heteroskedasticity has the potential to be very useful in practice.

[1] While the identification argument is in principle non-parametric, based simply on a path of variances, this path can only be recovered from the data by an econometrician under functional forms like GARCH. These moments are thus not available to the econometrician, in the sense of being consistently estimable, without strong non-parametric assumptions. These distinctions are discussed in further detail in Section 2.4.

I present a new identification argument based on heteroskedasticity that does not refer to the variance path, and thus need not make use of a particular parametric model. If time-varying volatility is present in any (unspecified) form, identification follows from the autocovariance of the volatility process. Since shocks are assumed to be uncorrelated over time, the autocovariance of squared residuals picks up only dynamics of the volatility process. This autocovariance furnishes equations that identify the response matrix and the structural shocks (up to an ordering) under very general conditions. In a simple model, the use of the autocovariance for identification can be motivated as an instrumental variables problem. The argument is similar in spirit to identification based on non-Gaussianity (e.g., Gouriéroux & Monfort (2015, 2017) and Hyvärinen, Zhang, Shimizu, & Hoyer (2010)) which uses different higher moments, but assumes that any volatility processes of the shocks are independent. I additionally derive testable implications of the identification condition, allowing it to be directly tested. Testing identification conditions has otherwise proven difficult when identification is based on heteroskedasticity, since the conditions typically apply to parameters only identifiable *conditional* on identification holding.

Identification based on time-varying volatility (TVV-ID) establishes identification under general conditions. Indeed, it separately establishes identification via a novel channel for the models that have previously been shown to offer identification via heteroskedasticity (e.g., GARCH of Sentana & Fiorentini (2001) and regimes of Rigobon (2003)). More importantly, it gives researchers the freedom to develop new alternative models and procedures in contexts exhibiting time-varying volatility, without having to stop to establish identification from scratch. As opposed to identification via heteroskedasticity being a model-dependent argument, TVV-ID progresses towards a model-free argument, which researchers can apply in ways that best suit their data. Importantly, it admits more flexible models where the volatilities are state variables as opposed to parameters, as in the stochastic volatility (SV) model. It is unique in demonstrating that the parameters of interest can be consistently estimated in such contexts, since identification follows from moments that can be consistently estimated even when the volatilities cannot.

These results mean that any estimator that fits an autocovariance to the squared residuals can implement TVV-ID. The most natural candidate is GMM, which needs no parametric assumptions. However, a researcher can also use a (quasi-) likelihood based on any model that implies such an autocovariance. I compare a variety of approaches considered in the literature with some newly-admissible estimators based on TVV-ID. I find that an estimator based on an AR(1) SV model performs best across many DGPs.

Identification via heteroskedasticity has been widely adopted in practice. Its use has spread from macrofinance to fields including public finance, growth, trade, political economy, agriculture, energy, education, marketing, and even fertility. This proliferation illustrates that there is potential value in freeing applied researchers from the strict parametric models they have been required to use and understanding any limitations of such models. The full flexibility offered by TVV-ID also shows that macro models including time-varying volatility are often estimated without realizing and exploiting its implications for identification. For example, Primiceri (2005) assumes a triangular H matrix when his volatility model means that these restrictions are not required for identification (and can thus be tested as overidentifying restrictions).

As an empirical application, I use TVV-ID to estimate fiscal multipliers and test previous identifying assumptions from the literature.[2,3] The multipliers I estimate are lower than those of Blanchard & Perotti (2002) or the comparative study of Mertens & Ravn (2014). I show that the narrative tax shocks often used for identification may not pass standard tests for validity. I reject the key parameter, the elasticity of tax revenues to output, obtained by both and obtain a value, 1.58, in line with Follette & Lutz's (2010) estimate based on institutional data. My multipliers accord with recent estimates of Caldara & Kamps (2017) and Ramey & Zubairy (2018).

The remainder of this paper proceeds as follows. Section 2 describes the identification problem in detail and presents the theoretical results. Section 3 compares implementations of TVV-ID and other identification schemes in simulation. The empirical application follows in Section 4. Section 5 concludes.

Notation

The following potentially unfamiliar notation is used in the paper. $\otimes$ represents the Kronecker product of two matrices; $\odot$ represents the element-wise product of two matrices (i.e. Hadamard product); $A_{(i)}$ denotes the i^{th} row of matrix A; $A^{(j)}$ denotes the j^{th} column of matrix A; A_{ij} denotes the ij^{th} element of matrix A; $A^{(-i)}$ denotes all columns of A except for the i^{th}, and similarly for rows and elements; $matdiag\left(A\right)$ is a vector of the diagonal elements

[2] This application provides an interesting test case as competing identification assumptions have been proposed for the same simple 3-variable reduced form model, which can be directly tested using TVV-ID. Further, no existing research has sought to exploit time-varying volatility in this setting, presenting an opportunity to develop a novel channel of identification.

[3] I have considered numerous other empirical applications. To summarize key results, I find that the recursive structure of Bernanke, Boivin, & Eliasz (2005) can be rejected, causing the price puzzle to return, and promoting other surprising behaviour at the contemporaneous horizon; the recursive structure of Kilian (2009) summarizes the data well (assumed zeros correspond to precisely estimated zeros); in Kilian & Park (2009) the zeroes assumed in the asset column of the contemporaneous response matrix are at odds with point estimates, but cannot be rejected; the assumptions of Blanchard & Quah (1989) are borne out strongly by TVV-ID; the exogeneity of uncertainty assumed in Bloom (2009) can be rejected, and the shapes of key responses to uncertainty shocks change somewhat.

of the square matrix A; $diag\left(a\right)$ is a diagonal matrix with the vector a on the diagonal; $x_{1:t}$ denotes $\{x_1, x_2, \ldots, x_t\}$.

Additionally, I use the non-standard notation $E_t\left[\cdot\right]$ to denote a time-specific expectation, i.e. the mean value of x_t at time t, as opposed to across t, and similarly $E_{t,s}\left[\cdot\right]$ when both time t, s variables are contained in the argument. This notation is used to make explicit that stationarity is not being assumed, unless otherwise noted, and to avoid the ambiguity (and possible non-existence) present in simply writing $E\left[x_t\right]$ in a non-stationary context. The use of E_t should not be confused with reference to the t information set; when a specific information set is intended, I condition on it explicitly.

2 Identification theory

In the canonical SVAR setting, a vector of innovations, η_t, is composed of unobserved structural shocks, ε_t, via a response matrix, H. This represents a more general decomposition problem. η_t is $n \times 1$, obtained from a reduced-form model or directly observed. For example, a structural vector auto-regression (SVAR) based on data Y_t would yield $A\left(L\right)Y_t = \eta_t$. Similarly, ε_t is $n \times 1$, so H is $n \times n$. Thus,

$$\eta_t = H\varepsilon_t, t = 1, \ldots, T, \tag{1}$$

leaving H and, equivalently, ε_t, to be identified. Equation (1) could also describe a factor model, for example. I begin by presenting a simple example under special assumptions to outline the identification problem and how heteroskedasticity may solve it. I then derive a representation of higher moments of the reduced-form innovations to serve as identifying equations. The following section establishes conditions under which these equations have a unique solution. I go on to highlight the role of the various assumptions and identification conditions, propose a simple test of the identification conditions, explain the relation to existing identification approaches, and discuss the interpretability of the shocks.

2.1 Intuition for the use of heteroskedasticity

Before the impact of heteroskedasticity can be illustrated, some standard assumptions underlying equation (1) are required.

Assumption 0. (temporary) *For all* $t = 1, 2, \ldots, T$,

1. $E\left[\varepsilon_t\varepsilon_t' \mid \sigma_t\right] = diag\left(\sigma_t^2\right) \equiv \Sigma_t$ *(σ_t^2 is the conditional variance of the shocks),*

2. σ_t *is a strictly positive stochastic process with time-invariant moments up to fourth order,*

3. $E[\Sigma_t] = \Sigma_\varepsilon$,

4. *Shocks satisfy conditional mean independence,* $E[\varepsilon_{it} \mid \varepsilon_{-is}] = 0$ *for all* i*, all* $t, s = 1, 2, \ldots T$,

5. H *is time-invariant, invertible, with a unit diagonal normalization.*

The fourth point substitutes conditional mean independence for the usual slightly weaker uncorrelated shocks assumption. While the variance of shocks may change, fixing H (as in Assumption 0.5) means that the economic impact of a unit shock remains the same. It is natural to seek to identify H from the overall covariance of η_t, $E[\eta_t \eta_t'] = \Sigma_\eta$. However, it is well-known that these equations can only identify H up to an orthogonal rotation, Φ ($\Phi\Phi' = I$).[4]

Variation in Σ_t may allow the researcher to overcome this indeterminacy. Consider a simple two-variable example, where one structural variance is time-varying and the other is fixed. This admits the simplest form of the Rigobon (2003) approach, which yields closed form solutions for H (as in Nakamura & Steinsson (2018), for example). Without loss of generality, assume σ_{2t}^2 changes and $\sigma_{1t}^2 \equiv \sigma_1^2$ is constant. Denote

$$\sigma_t^2 = \begin{bmatrix} \sigma_1^2 \\ \sigma_{2t}^2 \end{bmatrix}, \ H = \begin{bmatrix} 1 & H_{12} \\ H_{21} & 1 \end{bmatrix}.$$

The conditional variances of the reduced-form innovations are given by $E_t[\eta_t \eta_t' \mid \sigma_t] = H\Sigma_t H'$. Given two subsamples, A, B, containing the sets of time points T_A, T_B, it is shown by Rigobon & Sack (2004) (and in the Supplement) that

$$\frac{E_{T_A}[\eta_{1t}\eta_{2t}] - E_{T_B}[\eta_{1t}\eta_{2t}]}{E_{T_A}[\eta_{2t}^2] - E_{T_B}[\eta_{2t}^2]} = \frac{H_{12}\Delta(\sigma_{2t}^2)}{\Delta(\sigma_{2t}^2)} = H_{12}. \tag{2}$$

where the $\Delta(\cdot)$ operator represents the difference in expectation of the argument between subsamples T_A, T_B. Assuming that $\Delta(\sigma_{2t}^2) \neq 0$, H_{12} can thus be identified in closed form. σ_{2t} need only have finite second moments for all $t \in T_A, T_B$. While this argument is motivated by a regime-based process, identification holds even when misspecified, provided $\Delta(\sigma_{2t}^2) \neq 0$

[4] Observe $\Sigma_\eta = H\Sigma_\varepsilon H' = (H\Phi)(\Phi'\Sigma_\varepsilon\Phi)(H\Phi)' = H^*\Sigma_\varepsilon^* H^{*'}$, where $H^* = H\Phi D_{H,\Phi}$ and $\Sigma_\varepsilon^* = D_{H,\Phi}^{-1}\Phi'\Sigma_\varepsilon\Phi D_{H,\Phi}^{-1}$, with $D_{H,\Phi}$ the matrix that unit-normalizes the diagonal of $H\Phi$. This means that the pairs (H, Σ_ε) and $(H^*, \Sigma_\varepsilon^*)$ are observationally equivalent. Alternatively, note that due to the symmetry of Σ_η, it offers $n(n+1)/2$ equations, but there are n^2 unknowns. This is the fundamental identification problem posed by the SVAR methodology and indeed many related models (e.g., factor models).

and σ_1 is indeed fixed. If there are in fact regimes, they need not be known or correctly specified, as noted in Rigobon (2003). However, if the value of the σ_{2t} process is instead constant, $\Delta\left(\sigma_{2t}^2\right)$ would be zero in population, and identification fails.

Rigobon's approach exploits moment conditions based on subsample means of the variance process, but arguments are possible using other moments. Across periods, there is motivation for an instrumental variables (IV) approach. Noting

$$\begin{aligned} \eta_{2t}\eta_{1t} &= H_{21}\varepsilon_{1t}^2 + H_{12}\varepsilon_{2t}^2 + \varepsilon_{1t}\varepsilon_{2t} + H_{12}H_{21}\varepsilon_{1t}\varepsilon_{2t}, \\ \eta_{2t}^2 &= H_{21}^2\varepsilon_{1t}^2 + 2H_{21}\varepsilon_{1t}\varepsilon_{2t} + \varepsilon_{2t}^2, \end{aligned}$$

it is clear that H_{12} would be identified from the ratio of the $H_{12}\varepsilon_{2t}^2$ and ε_{2t}^2 terms. This is not possible as only the values of η_t are observed, and not their separate components. However, a lagged value of η_{2t}^2 can be used as an instrument for ε_{2t}^2. Note

$$cov\left(\eta_{2t}\eta_{1t}, \eta_{2(t-p)}^2\right) = H_{12}cov\left(\varepsilon_{2t}^2, \varepsilon_{2(t-p)}^2\right), \ cov\left(\eta_{2t}^2, \eta_{2(t-p)}^2\right) = cov\left(\varepsilon_{2t}^2, \varepsilon_{2(t-p)}^2\right),$$

by Assumption 0.4 and the fact that σ_1 is fixed. H_{12} is then identified in closed form:

$$\frac{cov\left(\eta_{2t}\eta_{1t}, \eta_{2(t-p)}^2\right)}{cov\left(\eta_{2t}^2, \eta_{2(t-p)}^2\right)} = \frac{H_{12}cov\left(\varepsilon_{2t}^2, \varepsilon_{2(t-p)}^2\right)}{cov\left(\varepsilon_{2t}^2, \varepsilon_{2(t-p)}^2\right)} = H_{12}. \tag{3}$$

This is the familiar IV estimator, where the dependent variable is $\eta_{2t}\eta_{1t}$, the endogenous regressor is η_{2t}^2, and the instrument is $\eta_{2(t-p)}^2$. This works because the previous value $\eta_{2(t-p)}^2$ is uncorrelated with all period t terms except those containing ε_{2t}^2. The argument applies for any lag, p. Identification holds provided

$$cov\left(\varepsilon_{2t}^2, \varepsilon_{2(t-p)}^2\right) \neq 0$$

for some p.

This requirement that the p^{th} autocovariance of η_{2t}^2 is non-zero is satisfied by a variety of processes for σ_{2t}^2. If the true process is regime-based, as suggested by the Rigobon estimator, identification follows from the non-zero autocovariance around break dates. In an SV model, it holds if the AR coefficient is non-zero. In a GARCH model at least one of the autoregressive parameters must be non-zero. This simple example displays the crux of TVV-ID: given the structure of the autocovariance of $\eta_t\eta_t'$, comparing elements of the autocovariance (in this simple case, via a ratio) identifies the columns of H.

This flexibility of identification – independent of specification – is not shared by the

Figure 1: Distribution of AR(1) coefficients of η_t^2

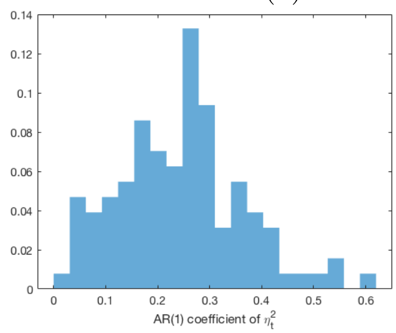

Time series η_t are obtained as reduced-form innovations from AR(12) processes fitted to each of McCracken & Ng's 128 FRED-MD monthly time series. The figure displays the distribution of the implied AR(1) coefficients of η_t^2.

existing approaches. I have made no assumptions about whether the heteroskedasticity is conditional or unconditional (either can imply a suitable autocovariance) and I have required only that the volatility process is stationary and exhibits some degree of persistence.

Empirically, there is strong evidence of such persistence, as discussed in Jurado, Ludvigson, & Ng (2015), for example. As a simple exercise, Figure 1 displays AR(1) parameters of η_t^2, where η_t are residuals of AR(12) models fitted to each series of McCracken & Ng's FRED-MD database in turn. I reject the null hypothesis of zero autocovariance at the 1% level for 96 of the 128 series, 5% for 98, and 10% for 101. A Ljung-Box test, as in Lanne & Saikkonen (2007), rejects homoskedasticity at the 1% level for 100 of the series and the 5% level for 103. The identifying condition is frequently satisfied empirically.

In this simple case, multiple autocovariances can easily be combined; each yields moments of the form

$$cov\left(\eta_{2t}\eta_{1t}, \eta_{2t-p}^2\right) - H_{12}cov\left(\eta_{2t}^2, \eta_{2t-p}^2\right) = 0,$$

which can be stacked to yield an overidentified GMM problem. Alternatively, it might be natural to assume that the (log) variances follow some loose parametric form, like an AR(1), and let this imply a whole range of autocovariances.

2.2 Identification via time-varying volatility

In the previous section, I made strong assumptions to assist intuition. I now relax them and develop TVV-ID in its general form. Again, let

$$\eta_t = H\varepsilon_t,\ t = 1, 2, \dots T.$$

Write $\mathcal{F}_{t-1} = \{\varepsilon_1, \dots \varepsilon_{t-1}, \sigma_1, \dots \sigma_{t-1}\}$. I replace Assumption 0 with Assumption A:

Assumption A. For every $t = 1, 2, \ldots, T$,

1. $E_t(\varepsilon_t \mid \sigma_t, \mathcal{F}_{t-1}) = 0$ *and* $Var_t(\varepsilon_t \mid \sigma_t, \mathcal{F}_{t-1}) = \Sigma_t$,
2. $\Sigma_t = diag(\sigma_t^2), \sigma_t^2 = \sigma_t \odot \sigma_t$,
3. $E_t[\sigma_t^2] < \infty$.

By explicitly conditioning on σ_t, these assumptions cover both SV and auto-regressive conditional heteroskedasticity-type (ARCH) models (where σ_t is a function of $\varepsilon_1, \ldots \varepsilon_{t-1}$), amongst many others, including unconditional heteroskedasticity. While this means that first-order shocks can have impacts on the second-order variance process, this setup does generally rule out first-order effects of innovations to volatility, as is also the case in existing approaches to identification exploiting time-varying volatility (and indeed VAR models more broadly). An exception arises when any first-order effects are driven by an observable factor of the volatility process, as in the model of Carriero, Clark, & Marcellino (2017).

In addition, I make an assumption on H:

Assumption B. *H is time-invariant, full rank, and has a unit diagonal.*[5]

The assumption that H is time-invariant is crucial for identification, ubiquitous in the literature. In fact, identification has not been established when H is time-varying except in very special cases, as in Angelini, Bacchiocchi, Caggiano, & Fanelli (2018). Work allowing more flexible time-variation in H is limited to Bayesian frameworks, most notably Cogley & Sargent (2005) or Primiceri (2005). I discuss the time-invariance of H further following the main identification results. Implicit in this setting and most related work is the additional assumption of invertibility, so the shocks are fundamental and thus recoverable from η_t.

Decomposition of $\eta_t \eta_t'$

To obtain moments in terms of just H and the underlying volatility process, I work with a transformation of η_t, (ζ_t, defined below), as my basic data. I begin by writing the decomposition

$$\eta_t \eta_t' = H \Sigma_t H' + V_t, \; V_t = H\left(\varepsilon_t \varepsilon_t' - \Sigma_t\right) H',$$

[5] The unit diagonal assumption is a normalization, without loss of generality. Note that even if there are zeros in H, such that certain column orderings are incompatible with a unit-diagonal, this poses no problem for the *identification* of H, since column order is imposed only ex-post for interpretation.

where Σ_t is unknown. Define L to be an elimination matrix, and G a selection matrix (of ones and zeros), see Magnus & Neudecker (1980), for example.[6] Then

$$\begin{aligned}\zeta_t &= vech\left(\eta_t \eta_t'\right) = vech\left(H\Sigma_t H'\right) + vech\left(V_t\right) \\ &= L\left(H \otimes H\right) vec\left(\Sigma_t\right) + v_t,\ v_t = vech\left(V_t\right) \qquad (4) \\ &= L\left(H \otimes H\right) G\sigma_t^2 + v_t, \qquad (5)\end{aligned}$$

The simplification from (4) to (5) in the first term is stark and follows due to the diagonality of Σ_t using A.2. From the definition of V_t, A.1, A.3, and B, $E_t\left[V_t \mid \sigma_t, \mathcal{F}_{t-1}\right] = 0$, so $E_t\left[v_t \mid \sigma_t, \mathcal{F}_{t-1}\right] = 0$ and

$$E_t\left[\zeta_t \mid \sigma_t, \mathcal{F}_{t-1}\right] = L\left(H \otimes H\right) G\sigma_t^2.$$

This provides a signal-noise interpretation for the decomposition of the outer product $\eta_t\eta_t'$. It follows from A.3 that I can integrate over Σ_t to obtain $E_t\left[v_t \mid \mathcal{F}_{t-1}\right] = 0$ and similarly that $E_t\left[|v_t|\right] < \infty$. Therefore v_t is a martingale difference sequence. Each observation of ζ_t $(\eta_t\eta_t')$ is an observation of $H\Sigma_t H'$, plus mean-zero noise.

Properties of ζ_t

Assumption C expands on A.3 to permit a characterization of the autocovariance of ζ_t.

Assumption C. For every t,

1. $Var_t\left(\sigma_t^2\right) < \infty$,

2. $Var_t\left(\varepsilon_t\varepsilon_t'\right) < \infty$.

Using these additional assumptions, the autocovariance of ζ_t has a convenient form:

Proposition 1. *Under Assumptions A.1-2, B, & C,*

$$Cov_{t,s}\left(\zeta_t, \zeta_s\right) = L\left(H \otimes H\right) GM_{t,s}\left(H \otimes H\right)' L',\ t > s \qquad (6)$$

where

$$M_{t,s} = E_{t,s}\left[\sigma_t^2 \sigma_s^{2'}\right] G' + E_{t,s}\left[\sigma_t^2 vec\left(\varepsilon_s\varepsilon_s' - \Sigma_s\right)'\right] - E_t\left[\sigma_t^2\right] E_s\left[\sigma_s^{2'}\right] G'.$$

[6] This means $vech\left(A\right) = Lvec\left(A\right)$ and $vec\left(ADA'\right) = \left(A \otimes A\right) Gd$ where $d = diag\left(D\right)$.

This equation represents an "observable" quantity, $cov_{t,s}(\zeta_t, \zeta_s)$, as a product of H and the $n \times n^2$ matrix $M_{t,s}$ (composed of $n \times (n^2+n)/2$ different moments of the underlying variance process). If $E_{t,s}\left[\sigma_{it}^2\left(\varepsilon_s\varepsilon_s' - \Sigma_s\right)\right]$ is diagonal (as in an SV model, or if any ARCH effects come from only *own* past shocks), $M_{t,s}$ can be replaced with $\tilde{M}_{t,s}G$ where $\tilde{M}_{t,s}$ is only $n \times n$.

An autocovariance of the vectorization of $\eta_t\eta_t'$ can thus be expressed as just a product of H, an $n \times n^2$ nuisance matrix, and known matrices of zeros and ones. This is remarkably parsimonious for a covariance of random matrices. Note that stationarity has not been assumed, merely the existence of higher moments. All of the expectations used are well-defined for an object at a particular point in time, even if the distribution might be different at another point in time. A single autocovariance provides $(n^2+n)/2 \times (n^2+n)/2$ equations in $n^2 - n + n(n^2+n)/2$ unknowns, so the order condition is satisfied.

Uniqueness

Having derived a set of equations of adequate order to identify H, it remains to show that they yield a unique solution. The conditions under which (6) yields a unique solution for H are established by Theorem 1.

Theorem 1. *Under Assumptions A.1-2, B, & C, equation (6) holds. Then H and $M_{t,s}$ are jointly uniquely determined from (6) (up to labeling of shocks) provided rank$(M_{t,s}) \geq 2$ and $M_{t,s}$ has no proportional rows.*

Theorem 1 states that (under certain conditions) equation (6) will yield a unique solution for the relative magnitudes of elements in each column of H. The identification result is based on period-specific moments – an autocovariance between two specific time periods, s, t – so stationarity is not assumed and is not required. In practice though, fourth-order stationarity of ε_t will often be needed so that (6) may be consistently estimated across the full sample. The solution is unique up to column order given the unit-diagonal normalization.[7] However, there are $n!$ column orderings. The same is true for any statistical identification approach, including those based on heteroskedasticity or non-Gaussianity, and is discussed in Chapter 14 of Kilian & Lütkepohl (2017). Ordering or otherwise labeling the columns, or equivalently the shocks, is an issue of rendering the shocks interpretable in an economic sense. In some cases, the labeling of shocks is unnecessary (as in factor models), and identification is complete, but for policy analysis labeling is required, as discussed in Section 2.5.

Theorem 1 makes two requirements of $M_{t,s}$. First, it must have rank of at least 2. Second, it must have no proportional rows. This is weaker than a full rank condition, since

[7] After a re-ordering of columns, H can be re-normalized to maintain a unit-diagonal.

rows of $M_{t,s}$ may be linear combinations, so long as they are not simply proportional. This dimensionality requirement ensures adequate heterogeneity in $M_{t,s}$ to uniquely identify H. $M_{t,s}$ encodes the autocovariance (and potentially ARCH) properties of σ_t^2 with σ_s^2; in an SV model, $M_{t,s}$ is simply the autocovariance of σ_t^2 multiplied by G'. These conditions imply that all n variances must be time-varying and persistent, and additionally that no two variance processes can have fully proportional autocovariance structures with respect to σ_s^2. Provided each variance has at least some persistent idiosyncratic component, this will not be the case. Jurado et al (2015) find that there are indeed strong idiosyncratic components in time-varying volatility that cannot be explained by common factors; the identification conditions will hold as long as those idiosyncratic components are persistent, and thus impact $M_{t,s}$. Conversely, in some finance settings (eg. Campbell, Giglio, Polk, & Turley (2017)), many volatilities are modeled as proportional. In Section 2.3, I propose a formal test of the identification conditions to evaluate these cases in practice. If the proportional row condition on $M_{t,s}$ does in fact fail, partial identification is still possible, as established in Corollary 1.

Corollary 1. *Under Assumptions A.1-2, B, & C, equation (6) holds. Then $H^{(j)}$ is identified from (6) provided rank $(M_{t,s}) \geq 2$ and $M_{t,s}$ contains no rows proportional to row j.*

This shows that columns of H pertaining to shocks whose volatility processes do not have proportional autocovariance structures can still be identified.

The identification conditions in Theorem 1 can be loosened by exploiting additional identifying equations. If, for example, the (often highly informative) mean

$$E_t\left[\eta_t \eta_t'\right] = E_t\left[\zeta_t\right] \tag{7}$$

is considered, Theorem 1 can be supplanted by Theorem 2.

Theorem 2. *Under Assumptions A.1-2, B, & C, equation (6) holds. Then H is uniquely determined from (6) and (7) (up to labeling of shocks) provided $\begin{bmatrix} M_{t,s} & E_t\left[\sigma_t^2\right] \end{bmatrix}$ has rank of at least 2 and no proportional rows.*

Theorem 2 shows that, provided the covariance of η_t is also used as an identifying moment, a proportional row assumption must additionally relate $E_t\left[\sigma_t^2\right]$ to $M_{t,s}$ in order for identification to fail. Similar arguments can be made, adding in further observable moments, requiring any proportionality extend to a matrix with progressively more columns. Corollary 1 can also be applied to Theorem 2. A major implication of Theorem 2 is described in Corollary 2.

Corollary 2. *H is uniquely determined from (6) and (7) (up to labeling of shocks) if at least $n-1$ shocks display time-varying volatility with non-zero autocovariance, provided that for no two shocks i, j, $cov_{t,s}\left(\sigma_{it}^2, vec\left(\varepsilon_s\varepsilon_s'\right)\right) = cov_{t,s}\left(\sigma_{jt}^2, vec\left(\varepsilon_s\varepsilon_s'\right)\right)\frac{E_t\left[\sigma_{it}^2\right]}{E_t\left[\sigma_{jt}^2\right]}$.*

Corollary 2 states that with the addition of (7), only $n-1$ dimensions of persistent time-varying volatility are sufficient to identify H, except in a very special case. This degenerate case amounts to the autocovariance structure of two shock variances being proportional by the ratio of their means. As discussed in Section 2.4, this weaker dimensionality requirement puts TVV-ID on a level footing with existing heteroskedasticity-based approaches.

Time-invariance of H

While TVV-ID focuses on the instability of the variances of structural shocks, H is assumed fixed. Although this is in principle a strong assumption, no existing identification scheme can flexibly accommodate time-varying H (Carreiro, Clark & Marcellino (2017) and Angelini et al (2018) do so under very specific functional forms). Even the simplest recursive short-run restrictions, when the true structure is in fact recursive, do not identify a known moment of H if H is in fact time-varying. Allowing H to vary more generally presents an interesting econometric problem, which warrants further study. While there are workhorse models in macroeconomics that allow for time-varying H, (e.g., Primiceri (2005) and Cogley & Sargent (2005)), these all adopt a Bayesian framework without identification results to separate variation in H from variation in Σ_t based on properties of the observable data alone. In this context, the parameter values obtained are driven by the structure of the priors, imposing information the data could never offer. As such, these approaches are largely orthogonal to the goal of this paper to provide non-parametric frequentist identification results facilitating consistent estimation of H based on observable data and (relatively) mild assumptions. While some frequentist work has adopted time-varying parameters (TVP) in the reduced form model, for example Auerbach & Gorodnichenko (2012), such papers are still unable to incorporate variation in H; time-variation in reduced form parameters can be combined with TVV-ID.

There are two ways in which time-variation in H is potentially compatible with TVV-ID. First, if H varies at a slower rate than the variances, identification may still hold asymptotically; H will be locally stationary over intervals over which the variances are not. Such a case could be explored in an infill-asymptotic setting, for example. Theoretical work sometimes reflects such distinctions in the rate of variation; for example, Barro & Liao (2017) split volatility into short-run and long-run components, with agents' behaviour driven by the slower moving component. Second, compared to identification exploiting regimes, as in Rigobon (2003), TVV-ID is better equipped to permit estimation over sub-samples over which H may plausibly be fixed, since the data do not already need to be subdivided for identification under a constant H. Should a researcher remain worried about the assumption of a fixed H, tests of overidentifying restrictions remain an option, as H is always over-identified by TVV-ID. Further, Andrews (1993) develops tests for parameter instability in a GMM context, for example the sup-Wald test, the conditions for which are satisfied for a variety of time-varying volatility models.[8]

2.3 Testing the identification conditions

Testing conditions for identification based on heteroskedasticity is difficult in general. The requirements for identification impose conditions on parameters that are only identified *conditional* on identification holding. In Sentana & Fiorentini (2001), the time paths of structural variances are required to be linearly independent, and in Rigobon (2003) the two (or more) sets of structural variances must be non-proportional. In TVV-ID, $M_{t,s}$ must have rank of at least 2 and no proportional rows. Given knowledge of the structural parameters, these conditions could easily be tested, but those parameters cannot be recovered without assuming identification. However, in Proposition 2, I derive testable implications of $M_{t,s}$ being full rank that pertain to the reduced form covariance $cov_{t,s}(\zeta_t, \zeta_s)$.

Proposition 2. *By construction,* $rank(M_{t,s}) = rank(cov_{t,s}(\zeta_t, \zeta_s)) = r$*; if* $r = n$*,* $M_{t,s}$ *is full rank and the identification conditions of Theorem 1 are satisfied.*

The implication of $r = n$ for $M_{t,s}$ is in general actually stronger than the condition required for identification, which requires only a rank of 2, with no proportional rows (rows that are not proportional but are otherwise linear combinations lower the rank of $M_{t,s}$ but not prevent identification). Thus, this condition $rank(cov_{t,s}(\zeta_t, \zeta_s)) = n$ can be viewed as conservative with respect to the true identification conditions for TVV-ID.

The problem of testing for identification is now reduced to testing the rank of the relevant autocovariance of ζ_t. Tests of matrix rank have been studied extensively, for example by Cragg & Donald (1996). For this purpose, I impose the assumption of fourth-order stationarity on ε_t (as discussed above), so that the matrix $cov(\zeta_t, \zeta_{t-p})$ can be consistently estimated. Then, Theorem 3 provides a test statistic and asymptotic distribution to assess

[8]The less-familiar assumptions needed in Andrews (1993), those of Near-Epoch Dependence (NED), can be replaced by stronger properties that hold for both GARCH and SV processes. Lindner (2009) shows that GARCH satisfies β-mixing (and thus α-mixing with exponential rate) and Davis & Mikosch (2009) show that SV models inherit the mixing properties of the log-variance process. Andrews' (1983) results show that an AR(1) variance process is $\alpha-$mixing with exponential rate. These mixing properties can be shown to imply NED; see Davidson (1994) Chapter 17 for additional background.

the rank of the autocovariance matrix, and thus test whether the conditions to identify H using TVV-ID hold.

Theorem 3. *If* $\widehat{cov(\zeta_t, \zeta_{t-p})}$ *is an asymptotically normal estimator of* $cov(\zeta_t, \zeta_{t-p})$*, then under the null hypothesis that the autocovariance has rank* r*, the associated Cragg-Donald statistic* $CD_{\zeta,p}(r)$ *has the asymptotic distribution* $CD_{\zeta,p}(r) \xrightarrow{d} \chi^2\left(\left((n^2+n)/2 - r\right)^2\right)$.

The interested reader should consult Cragg & Donald (1996) for additional technical details and a description of the test statistic. Essentially, the test assesses the deviation of part of the estimated matrix from zero following r steps of Gaussian elimination. Note that since $cov(\zeta_t, \zeta_{t-p})$ is of dimension $(n^2+n)/2 \times (n^2+n)/2$, its rank will in general be greater than n in finite samples; indeed, Cragg & Donald's Assumption 1 requires this to be the case in finite samples.[9] While this result offers an immediate way to test the identification condition, as always it is unclear what constitutes a suitable level for a test of *strong* identification. I sketch general methods for assessing weak identification in nonlinear models, which could be applied here, in the Supplement.

2.4 Relation to existing approaches

TVV-ID generalizes the conditions under which previous approaches have established identification via heteroskedasticity and nests the parametric models on which they have relied. Below, I describe the relation of TVV-ID to each of the existing identification results.

Sentana & Fiorentini (2001) offer an identification argument that is in principle nonparametric; they show that, conditional on the time path of reduced form covariances, $\Sigma_{\eta 1:T}$, H is identified, provided the variance processes are linearly independent. However, this path is not in general available to the econometrician, who observes only the noisy $\eta_t \eta_t'$ in each time period, no matter the sample length. This leads the authors to recommend a GARCH functional form, which is unique in allowing the reduced form covariances to be deterministically recovered from the observations conditional on H and the parameters of the volatility process. Of particular concern is the fact that the need to recover the path of reduced form covariances to use as identifying moments rules out all variance processes where the variance innovations are not coupled to innovations to observable variables. This precludes all variance processes including state variables, in particular the very popular class of stochastic volatility models. The only alternative is to collapse observations into subsamples and apply the identification argument to the covariance path across these subsamples,

[9] Tests for rank exceeding n provide a possible avenue for a test of misspecification, which I defer to future work. The possibility that a test for identification may interact with evidence of misspecification is a general concern and not limited to the present setting.

essentially the Rigobon (2003) argument. However, if the true variance process is stationary, then, asymptotically, such an implementation will recover the same covariance across each subsample, and Sentana & Fiorentini's (2001) identification condition will fail.

TVV-ID avoids these issues entirely by not making any reference to the variance path for identification, instead using a single unconditional moment, the autocovariance of ζ_t. Because it is unnecessary to recover the variance path for identification, TVV-ID can admit a near arbitrary range of volatility models, and is truly non-parametric. Such moments can, under suitable assumptions, be consistently estimated even in models with state variables. TVV-ID is the first such scheme to imply that H can be consistently estimated *even when the volatility path cannot.*

Additionally, TVV-ID nests the implementations of Sentana & Fiorentini (2001) that have appeared in the literature, exclusively based on GARCH volatility processes. This is because a (stationary) GARCH process clearly implies a suitable matrix $M_{t,t-p}$ for autocovariance p.[10] Sentana & Fiorentini (2001) require $n-1$ dimensions of linearly independent time-varying volatility; TVV-ID similarly requires $n-1$ volatility processes with non-proportional autocovariance structures (Corollary 2). These conditions will generally coincide.

In a recent paper, Bertsche & Braun (2018) use the Sentana & Fiorentini (2001) argument to motivate identification of an SVAR based on SV. While a good heuristic argument, identification does not hold in a true sense since the moments on which the identification argument is based, $\Sigma_{\eta 1:T}$, cannot be recovered *even asymptotically* in a state space model like SV. There is no guarantee that the noise unavoidable in the versions of these moments available to the econometrician will not confound identification. Likewise, since the identifying moments $\Sigma_{\eta 1:T}$ cannot be consistently estimated, H cannot be argued to be consistently estimated. On the basis of TVV-ID though, their model is clearly identified from the unconditional moments in (6), which are consistently estimated (indirectly) via their EM approach. This means that their results are all valid, just on the basis of TVV-ID, which operates in the background.

As noted above, the Rigobon (2003) argument is essentially the same as the Sentana & Fiorentini (2001) argument, except that periods are pooled into subsamples for identification. $n-1$ shocks must exhibit non-proportional variance changes across regimes. The subsamples can either be based on external information (e.g., monetary policy announcement days) or estimated. The former is ideal, but puts an additional informational burden on the econometrician to supply the information. I show in Section 5 of the Supplement that the latter process of estimating regimes can induce bias in estimates, since the estimated regimes

[10] Milunovich & Yang (2013) offer an additional (local) identification argument for the GARCH model based on reduced-form moments, more similar to the TVV-ID approach.

may be endogenous with respect to the structural shocks. A third alternative of arbitrary regular cuts in the data (e.g split at $T/2$) causes identification to break down asymptotically if the volatility process is stationary (the moments converge over each subsample). Markov switching (e.g., Lanne, Lütkepohl, & Maciejowska (2010)) and smooth transition (e.g., Lütkepohl & Netšunajev (2017)) variants of the Rigobon argument address the regime estimation problem internally through their likelihoods. Because TVV-ID relies on full-sample unconditional moments, these challenges are avoided entirely. If the true volatility model is a regime-based model, TVV-ID can also nest that parametric form. While, within a regime, the volatilities are invariant, so there is no autocovariance, transitions between regimes induce autocovariance in the volatilities, providing the identifying variation needed by TVV-ID.[11]

TVV-ID belongs more broadly to a long literature (dating to at least Darmois (1953) and Skitovich (1953)) of identification based on higher moments. Work has generally focused on contemporaneous moments (or cumulants), whereas TVV-ID exploits inter-temporal moments. This literature includes extensive work by Hyvärinen and co-authors (e.g., Hyvärinen, Karhunen, & Oja (2001)), which considers the model as a signal extraction problem, developing variants of the Independent Components Analysis approach to exploit non-Gaussianity. Provided $n-1$ shocks exhibit unconditional non-Gaussianity, H can be identified. Identification via non-Gaussianity is growing in prominence in economics (e.g., Gouriéroux & Monfort (2015, 2017)). In principle, non-Gaussianity encompasses heteroskedasticity, as time-varying volatility makes Gaussian shocks unconditionally non-Gaussian. TVV-ID cannot nest these highly general identification results. However, the converse is also true: identification via non-Gaussianity requires that, for $i \neq j$, shocks ε_{it} and ε_{jt} be mutually independent, not just orthogonal. This rules out dependence in higher moments, and thus restricts any volatility processes to be uncorrelated across shocks. Such an assumption is at odds with many empirical findings suggesting factor structures in macroeconomic volatilities, see Jurado et al (2015), for example. This means that there may be value to identification approaches tailored to heteroskedasticity when it is time-varying volatility that motivates the presence of non-Gaussianity, as borne out in the simulation study in Section 3.

The identification conditions for TVV-ID parallel those of these other approaches. Sentana & Fiorentini (2001), Rigobon (2003), and identification via non-Gaussianity all impose a dimensionality condition to ensure adequate heterogeneity to identify H, requiring $n-1$ non-proportional dimensions of time-varying volatility or $n-1$ shocks to be non-Gaussian.

[11] As a simple example, consider a univariate process with $\sigma_t^2 = 1, t = 1, \ldots, T/2$ and $\sigma_t^2 = 2, t = T/2 + 1, \ldots, T$. $cov\left(\sigma_t^2, \sigma_{t-1}^2\right) = \frac{1}{T}\left(2 + 1 \times (T/2 - 1) + 4 \times (T/2 - 1)\right) - \left(\frac{1+2}{2}\right)^2$, which converges to $2.5 - 1.5^2 = 0.25$ as T goes to infinity, so even with a single regime switch, the autocovariance is non-zero asymptotically.

Likewise, $n-1$ dimensions of time varying volatility with non-proportional autocovariance structures will satisfy the conditions for TVV-ID. However, TVV-ID does not require the researcher to recover paths for covariances and is thus able to accommodate a much wider range of volatility processes due to its non-parametric nature.

A final strength of TVV-ID is the testability of the identification conditions, as demonstrated in Theorem 3. This test exploits evidence of the identification conditions that can be found in the reduced form moment $cov(\zeta_t, \zeta_{t-p})$. This result is significant. The parameters to which identification conditions based on heteroskedasticity apply cannot be recovered without assuming identification holds. This means that such conditions cannot generally be tested directly. This poses difficulty in the GARCH implementation of Sentana & Fiorentini (2001). Lanne & Saikkonen (2007) propose a test for the dimensionality of a variance process based on time series of the structural shocks. However, the authors must obtain estimated time series of the structural shocks to test. As a result, they recommend assuming a recursive structure for part of H when the null hypothesis specifies only enough dimensions of heteroskedasticity to partially identify H, in order to avoid assuming heteroskedasticity. This means that frequently their test may be for heteroskedasticity in a *rotation* of the structural shocks of interest.[12] Lewis (2018) proposes tests for weak identification (and thus identification more broadly) in the Rigobon (2003) model, but the tests remain computationally challenging; recent work by Lütkepohl et al (2018) promises to offer an alternative closer in spirit to the test presented here.

2.5 Interpreting results

Having identified the columns of H through TVV-ID, it is frequently still necessary to label the columns of H, or, equivalently, the resulting structural shocks. Kilian and Lütkepohl (2017) discuss how there may in fact be some difficulty in interpreting these as economically meaningful shocks, given the purely statistical methods used to derive them; this step helps to develop such interpretations. In the Supplement, I outline a number of potential approaches to labeling the columns of H. These frequently constitute a weaker version of standard structural identification assumptions that might otherwise be used to identify the model. For example, instead of assuming a certain macroeconomic shock has no contemporaneous impact on some series (as in a Cholesky ordering, say), the recovered shock that has the closest-to-zero effect on that series could be labeled as that particular macroeconomic shock.

Any such labeling exercise does not, however, necessarily assume the shocks are meaning-

[12] While an orthogonal rotation of the variance paths themselves would not alter the dimensionality of the variance process, the test proposed is on the autocovariance of the recovered shock series, for which such a result is not obvious.

ful - it is possible that no shock meets a theoretically-motivated labeling criterion satisfactorily. A researcher so concerned can test whether a statistically-recovered shock represents a particular economic shock by formally testing conventional identifying assumptions as overidentifying restrictions. An alternative is to informally evaluate the extent to which the impulse response functions (IRFs) align with those based on economic theory, as in Brunnermeier et al (2017) or Lütkepohl & Netšunajev (2014).

Importantly, inference approaches that are valid for an estimated $\widehat{H}$ will also be valid for a labeled column of $\widehat{H}$, denoted $\widehat{H^{(j)}}$, under standard conditions. In general, the use of statistical measures to select a column of an estimated matrix will impact the asymptotic distribution of the ultimate column estimates. However, for most statistical labeling criteria that select a unique shock, the labeling criterion is consistent in the probability limit sense. This means that as $T \to \infty$, the probability of selecting the correct column based on the criterion approaches unity. Pötscher (1991) establishes asymptotic distributions in a discrete model selection setting building on intuition dating back to at least Geweke & Meese (1981). For a consistent labeling criterion, it is direct to show that a strong form of Pötscher's results hold. This means that if a labeling method is consistent and the asymptotic distribution of $\widehat{H}$ is known, the selected column $\widehat{H^{(j)}}$ simply inherits that asymptotic distribution. In other words, the labeling problem can be ignored for the purpose of asymptotic inference.

To counterbalance the challenge of shock labeling, an advantage of statistical approaches to identification is that it is straightforward to describe the impact of economic assumptions quantitatively when they are used. Because they are used to make discrete decisions – one shock or another is the policy shock – it is possible to report and compare estimated economic effects under alternative labeling assumptions. Frequently in empirical applications a large number of labeling assumptions will agree on the policy shock. Reporting these findings makes a single result compelling to readers believing any, but perhaps not all, of that collection of assumptions.

3 Estimators and performance

A strength of TVV-ID is that it is an identification argument not tied to any model or estimator. It can thus be implemented by any estimator that fits an autocovariance of the residuals to the data. This can either be explicit – in the case of GMM on equation (6) – or implicit, in the case of many likelihood models. This is in contrast to the Sentana & Fiorentini or Rigobon arguments, which require either a GARCH-type model or regimes, respectively. This means that a researcher can choose a completely non-parametric approach (GMM), whatever model she thinks best describes the data (quasi-maximum likelihood, QML), or compare a variety of different models for robustness.

As noted briefly above, while identification does not require stationarity, in order for all of the estimators considered here to be well-behaved, some degree of stationarity must be assumed. In the case of GMM, for example, ε_t must be fourth-order stationary so that the identifying moments (6) can be consistently estimated. In an SV model, $\log \sigma_t^2$ must be second-order stationary so the parameters of the SV process governing the autocovariance can be consistently estimated.

While GMM is the natural entirely non-parametric implementation of TVV-ID, the higher moments used for identification can be very noisily estimated in realistically short macro time series. This motivates the use of likelihood approaches, which make parametric assumptions in exchange for possible efficiency gains. QML is a natural way to incorporate the identifying information of multiple autocovariances implied by a functional form. The drawback of any likelihood-based approach is the necessity of specifying a law of motion for the structural variances. To some extent this may seem a return to parametric assumptions this paper set out to avoid. However, thanks to the general identification arguments offered above, *identification* is not tied to a particular functional form. In particular, the SV model is a common, highly flexible model of time-varying volatility that decouples the innovations in variances from the shocks themselves. It has proven popular in the financial econometrics literature, where much work has compared its ability to describe the data with GARCH and other models (e.g., Diebold & Lopez (1995), Kim, Shephard, & Chib (1998), Barndorff-Nielsen & Shephard (2002)). There is reason to believe it could at least be a competitor to GARCH-based approaches. Bertsche & Braun (2018) adopt the model to estimate an SVAR under heteroskedasticity (without the theoretical justification offered by TVV-ID), and find it performs well in simulation. Carriero, Clark, & Marcellino (2018) use it to capture time-varying volatility in an SVAR (identification follows from particular model features), as do many Bayesian applications (e.g., Uhlig (1997), Cogley & Sargent (2005), Primiceri (2005)). In this section, I put the three heteroskedasticity-based identification schemes discussed in this paper, as well as identification based on non-Gaussianity, to the test. I consider several different implementations in a simulation study based on a wide range of DGPs.

3.1 DGPs and estimators

The DGPs are empirically calibrated from the residuals of a bivariate SVAR where the two variables are the first principal component extracted from the McCracken & Ng FRED-MD database and the Fed Funds rate, identified using an AR(1) SV model. For each alternative volatility process considered, the model is calibrated based on the structural shock series resulting from the AR(1) SV estimates. The H matrix used in the simulations is

$$H = \begin{bmatrix} 1 & 0.298 \\ 0.033 & 1 \end{bmatrix}.$$

H_{21} represents the contemporaneous effect of a macroeconomic shock on the Fed Funds rate, and H_{12} represents the contemporaneous effect of a Fed Funds shock on the first principal component of the FRED-MD database. The DGPs thus calibrated consist of a Markov switching model, a GARCH(1,1) model (including a "weak" variant), and an AR(1) SV model (including varied sample size, a "weak" variant, and non-Gaussian (t_7) disturbances). I take 5000 replications, and unless otherwise noted, $T = 200$. Values of the parameters for the volatility models can be found in the Supplement.

I consider a diverse range of identification approaches and associated estimators. For TVV-ID, I consider an AR(1) SV QML implementation generalizing the EM algorithm of Bertsche & Braun (2018), exploiting the expansions of Chan & Grant (2016) in the E-step. I also use a 2-step GMM estimator, making use of the first autocovariance augmented by $E[\eta_t \eta_t']$. For the Sentana & Fiorentini approach (2001), I adopt two estimators based on the GARCH model. First, I estimate the standard GARCH(1,1) model adopted in Normandin & Phaneuf (2004), Lanne & Saikkonen (2007), Milunovich & Yang (2013), Lütkepohl & Milunovich (2016), and many others. Second, I consider a "hybrid GARCH" estimator, a GARCH(1,1) model where the autoregressive parameters are calibrated to macro data, but the mean parameters are estimated (details of this calibration based on the FRED-MD database can be found in the Supplement); this can be thought of as a parametric kernel estimator. For the Rigobon approach, I use estimated regimes based on the trace, median threshold, and 13 period windows, as recommended by results in Section 5 of the Supplement. I also use an arbitrary split at $T/2$, as well as a Markov switching model estimated via ML. Finally, for non-Gaussianity, I use the FastICA algorithm described in Shimizu, Hoyer, Hyvärinen, & Kerminen (2006) (preliminary simulations show performance superior to ML approaches). Details on selected estimators can be found in the Supplement. Columns of H are labeled using the infeasible method of minimizing the L^2 norm to the true value.

Table 1: Mean estimates and rejection rates

		QML AR(1) SV		GMM		Hybrid		GARCH		Sub-sample (rolling)		Sub-sample ($T/2$)		Markov Switching		Non-Gaussianity	
		mean	α	mean	α	mean	α	mean	α	mean	α	mean	α	mean	α	mean	α
Markov switching, $T = 200$	H_{21}	0.03	6.8	0.01	39.9	0.02	10.0	0.03	47.2	0.02	17.0	0.01	22.6	0.03	4.1	0.03	38.5
	H_{12}	0.28	10.1	0.45	44.2	0.34	12.2	0.31	45.4	0.38	4.3	0.38	4.2	0.28	4.5	0.27	43.7
	RMSE	2.70		6.78		5.49		4.87		4.45		6.61		2.45		4.86	
GARCH(1,1), $T = 200$	H_{21}	0.03	5.3	0.03	26.6	0.03	4.3	0.03	4.8	0.03	15.8	0.03	15.6	0.03	11.2	0.03	11.4
	H_{12}	0.29	6.8	0.40	32.0	0.33	5.3	0.30	4.7	0.19	2.3	0.36	2.8	0.37	11.1	0.41	13.3
	RMSE	2.96		7.73		2.47		2.58		5.47		6.89		5.28		6.98	
GARCH(1,1), $T = 200$, weak	H_{21}	0.03	19.0	0.02	48.7	0.02	24.4	0.03	4.8	0.03	14.1	0.02	21.6	0.03	9.6	0.02	8.1
	H_{12}	0.32	21.7	0.84	51.9	0.98	24.6	0.27	5.8	0.11	1.6	0.83	2.3	0.58	9.5	1.19	11.2
	RMSE	8.52		12.15		8.27		6.94		7.99		13.28		11.05		15.02	
AR(1), $T = 100$	H_{21}	0.03	14.9	0.01	45.2	0.03	9.2	0.03	22.3	0.02	19.4	0.02	17.0	0.02	11.5	0.02	31.2
	H_{12}	0.29	16.3	0.74	49.8	0.35	10.2	0.37	21.5	0.38	4.3	0.39	3.4	0.43	10.6	0.41	31.7
	RMSE	5.47		9.19		6.67		6.35		8.08		7.42		7.16		7.80	
AR(1), $T = 200$	H_{21}	0.03	7.8	0.02	40.6	0.03	6.9	0.03	23.4	0.02	19.1	0.02	18.9	0.03	5.4	0.03	35.4
	H_{12}	0.29	9.6	0.51	44.1	0.32	7.7	0.30	22.1	0.37	3.3	0.36	3.3	0.34	5.1	0.33	38.5
	RMSE	2.89		7.89		4.31		3.91		6.91		6.29		5.20		5.92	
AR(1), $T = 400$	H_{21}	0.03	4.2	0.02	33.7	0.03	4.3	0.06	61.7	0.02	25.2	0.02	19.9	0.03	5.5	0.03	37.1
	H_{12}	0.30	5.9	0.45	38.7	0.29	4.9	0.74	51.6	0.31	3.4	0.37	3.1	0.28	5.1	0.33	41.1
	RMSE	1.42		6.28		2.52		8.50		6.23		5.72		3.17		4.42	
AR(1), $T = 200$, weak	H_{21}	0.03	41.5	0.01	44.7	0.02	47.6	0.02	8.6	0.04	23.9	0.02	28.3	0.02	15.5	0.01	10.1
	H_{12}	0.33	42.9	0.63	45.2	0.50	48.6	0.50	10.1	0.09	2.8	0.47	3.6	0.51	14.9	0.52	12.0
	RMSE	7.63		9.08		8.36		7.80		6.38		8.41		8.00		8.96	
AR(1), $T = 200$, t_7 shocks	H_{21}	0.03	4.7	0.02	40.1	0.03	8.7	0.03	32.8	0.02	18.8	0.02	17.6	0.03	4.0	0.03	32.4
	H_{12}	0.30	5.8	0.60	44.0	0.32	9.6	0.31	30.7	0.33	3.6	0.38	3.6	0.32	3.5	0.30	34.4
	RMSE	2.24		8.45		5.62		4.95		6.38		6.27		5.65		4.54	

True values: $H_{21} = 0.033$, $H_{12} = 0.298$, nominal size $\alpha = 5\%$

Mean estimates for the full range of estimators for the specified DGPs. True $H_{21} = 0.033$ and $H_{12} = 0.298$. Labeling proceeds via an infeasible method matching H estimates to the true H to minimize L_2 norm. Rejection rates, α, are presented for a nominally-sized 5% test for each draw. Details on standard errors can be found in the Supplement. Since the RMSE must account for error in multiple parameter estimates, the MSE is computed for each, and then normalized by the square of the true parameter, before the root of the sum is taken.

3.2 Results

Table 1 reports the results. It lists the mean estimates for the off-diagonal elements of H, RMSE (root of weighted sum of MSEs for both parameters), and rejection rates for nominal 5% tests of the true parameter values using each estimator's appropriate standard errors (described in the Supplement). The choice of a MSE measure mirrors the related contemporary study of Lütkepohl & Schlaak (2018). Histograms reported in the Supplement show that distributions for most estimators and DGPs are centered around the true parameters; large discrepancies in mean estimates are mostly driven by outliers. Recall that the true values are $H_{21} = 0.033$ and $H_{12} = 0.298$. The former represents the contemporaneous response of the Fed Funds rate to a macroeconomic shock and the latter represents the contemporaneous response of the first principal component of the FRED-MD database to a Fed Funds shock.

Across DGPs, the QML implementation of the AR(1) SV model performs best. The mean estimates are accurate, and even when misspecified the RMSE is often only slightly worse than that for well-specified estimators. This makes it a compelling choice to implement TVV-ID. A further benefit is that tests of true values are fairly well-sized, except in the presence of weak identification.

The hybrid GARCH estimator and Markov switching estimators offer the next best performance. The mean estimates are still accurate, but their RMSEs are higher in general. They struggle in the face of weak variation in volatility. For the hybrid, this is largely because the calibrated parameters are no longer a good fit for the data. The standard errors for both estimators offer minimal size distortions, apart from cases of weak identification.

The FastICA estimator exploiting non-Gaussianity is also reliable. The mean estimates are close to the true values except for DGPs with small sample sizes or weak variation. In these cases, the higher moments on which this identification rests seem very imprecisely estimated – moreso than the *persistence* of the process, which TVV-ID exploits. In contrast, estimators like that for SV or GARCH models exploit a path of variances for identification as well as these unconditional higher moments of the data. The RMSE is accordingly higher, depending on the DGP. Naturally, its performance improves when disturbances are themselves non-Gaussian. The standard errors perform quite poorly with respect to rejection rates – this is because the asymptotic variance depends on up to the sixth moment of the shocks, so is very imprecisely estimated.

The GARCH estimator is generally competitive with the previous approaches, but breaks down for SV with $T = 400$. This is because the empirically calibrated DGP dictates parameters that are very close to non-stationarity when approximated by GARCH. As a result, with a longer draw of data, there is a reasonable chance of observing dynamics that appear explosive from a GARCH-fitting perspective, negatively impacting the estimates. This phenomenon also appears in un-reported simulations for different empirical calibrations, generally manifest in excess mass around zero for the H parameters when the GARCH parameters are close to the boundary of stationarity. Since these calibrations are empirical, this is a strike against adopting GARCH estimators for identifying SVARs in similar macro data. The rejection rates are accurate when well-specified, but as expected, break down when misspecified.

The Rigobon estimates based on rolling windows are quite good, which is unsurprising given the tuning parameters are optimized based on simulations discussed in Section 5 of the Supplement. Other combinations might harm performance. However, the breakdown is dramatic for weak identification; when changes in the volatilities are minimal, estimated regimes will be increasingly determined by realized shock values instead, biasing estimates, as discussed in detail in Section 5 of the Supplement. This estimator is not in general competitive with the best estimators in terms of RMSE. Generally, the same remarks apply to the simple $T/2$ split estimator. For both, the rejection rates are badly distorted.

GMM generally struggles, especially with small samples and weak variation. Since it relies mostly on fourth moments, for identification, without any auxiliary information from variance paths, this makes sense, as these moments are noisily estimated in those DGPs. Accordingly, the rejection rates are also distorted. For progressively higher T, additional simulations suggest performance does become acceptable. Thus, for larger sample sizes, GMM may offer a viable alternative requiring no parametric assumptions.

The results of this simulation study are related to those of Lütkepohl & Schlaak (2018). That paper estimates a range of parametric volatility models for a variety of DGPs and assesses the consequences of misspecification while comparing the performance of popular model-selection criteria with the goal of choosing a well-specified volatility model. They consider the MSEs for impulse response functions estimated using their competing models. Their results also demonstrate that a GARCH-based estimator can perform quite poorly under misspecification, with a striking break-down as T increases. The Markov switching estimator also performs reasonably well across DGPs considered in their study. Their study focuses on the single-break, smooth transition, Markov switching, and GARCH models, justified by existing identification results, and thus is not informative about the SV model found to be most reliable here.

4 Empirical application: fiscal multipliers

Considerable work has been devoted to estimating the value of fiscal multipliers, but has resulted in considerable disagreement over their size. The range of estimates is documented by Mertens & Ravn (2014), Caldara & Kamps (2017), and Ramey (2011a). Prominent estimates range from less than zero to over three. While government spending multipliers are perhaps most familiar, tax multipliers capture an equally important dimension of fiscal policy, and are central to current policy debates. Blanchard & Perotti (2002) (henceforth BP) is seminal in the literature; recent work by Mertens & Ravn (2013, 2014) and Mountford & Uhlig (2009) has obtained contrasting estimates. Caldara & Kamps (2017) show the discrepancy can be largely explained by differing values for the elasticity of tax revenues with respect to output. BP calibrate this parameter to 2.08 based on institutional information, Mountford & Uhlig's (2009) penalty-function identification is consistent with a prior for the elasticity centered around 3, and Mertens & Ravn (2014) (henceforth MR) estimate a value of 3.13 using Romer & Romer (2010) (henceforth RR) narrative shocks as external instruments. This setting provides an ideal test case for TVV-ID due to the relatively small dimension of the standard model ($n = 3$) and because TVV-ID offers a channel of identification completely different to those previously considered. This setting was not previously a strong candidate for identification via heteroskedasticity due to a lack of ex ante natural variance regimes.[13]

Theorem 1 shows that the autocovariance of volatility present in the data can identify the structural parameters determining fiscal multipliers without the economic assumptions required in prior work. BP need a calibrated value, the assumption of no contemporaneous response of spending to output, and a recursive ordering between tax revenue and spending, and MR require their instrument to be valid and there to be no contemporaneous response of spending to output. Since I am able to depart from these assumptions (making them over-identifying restrictions), I can test them using the results of TVV-ID.

4.1 Data & model

I adopt MR's trivariate VAR with federal tax revenue, federal government consumption and investment, and GDP, based on quarterly BLS data found in the NIPA tables, spanning 1950Q1 to 2006Q4.[14] Additional details on the data and de-trending procedures (including

[13] Higher dimensions put more strain on the EM algorithm required to implement TVV-ID using the favoured AR(1) SV model. While comparison with prominent papers on the effects of monetary policy that already incorporate time-varying volatility, like Primiceri (2005) or Brunnermeier et al (2017), is in principle interesting, these papers work with higher-dimensional models and are Bayesian in approach. This means that testing identification conditions is not a straightforward problem, and does not permit an "apples to apples" comparison, since "identification" follows at least partly from priors (particularly in Primiceri (2005)). In contrast, the fiscal multipliers debate provides competing identification assumptions based on the same simple reduced form model, which can be directly tested.

[14] While Auerbach & Gorodnichenko (2012) estimate a TVP model in the reduced form, I maintain the constant parameter reduced form VAR specification, since that remains the benchmark in the literature, even in more recent work, such as Caldara & Kamps (2017) and Mertens & Montiel Olea (2018). Maintaining the same reduced form allows a clearer comparison of identifying assumptions across approaches.

Figure 2: Moving averages of squared residuals and shocks

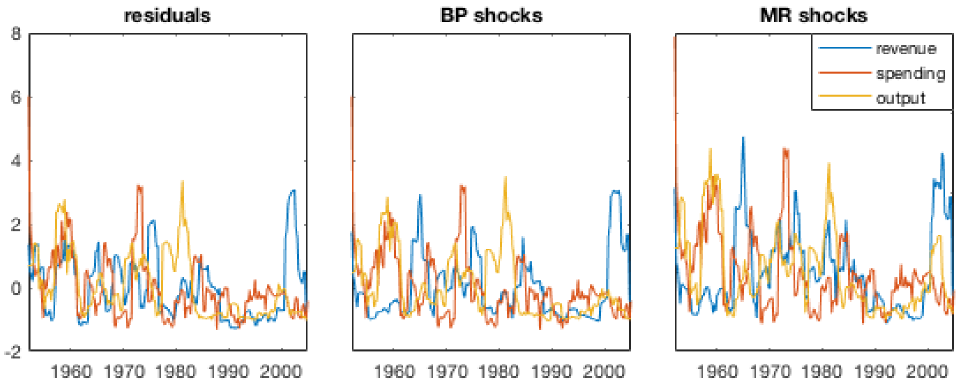

2-year moving averages of the square of the specified series. For the first panel, this is the reduced form residuals, BP structural shocks for the second, and MR structural shocks for the third.

Table 2: Tests of identification assumptions

null/alternative	$\widehat{cov\left(\zeta_t, \zeta_{t-1}\right)}$
$rank\left(\cdot\right) = 1/rank\left(\cdot\right) > 1$	112.24^{***}
$rank\left(\cdot\right) = 2/rank\left(\cdot\right) > 2$	253.70^{***}

Cragg-Donald (1996) tests of the rank of $\widehat{cov\left(\zeta_t, \zeta_{t-1}\right)}$, where $\zeta_t = vech\left(\eta_t \eta_t'\right)$; test statistics are starred at the 1% level and follow the $\chi^2\left(\left(\left(n^2+n\right)/2-r\right)^2\right)$ distribution. A White variance matrix is used; the rejection is stronger still using two unreported alternative HAC approaches.

federal vs. general government data) can be found in MR. I use the replication code available on Mertens' website to obtain identical residuals.

In MR's notation, the BP benchmark model is

$$\begin{aligned}
u_t^T &= \sigma_T e_t^T + \theta_G \sigma_G e_t^G + \theta_Y u_t^Y \\
u_t^G &= \gamma_T \sigma_T e_t^T + \sigma_G e_t^G + \gamma_Y u_t^Y \\
u_t^Y &= \zeta_T u_t^T + \zeta_G u_t^G + \sigma_Y e_t^Y,
\end{aligned}$$

where $u_t = \eta_t$ and e_t are structural shocks with $E\left[e_t^{'} e_t\right] = I$. Key parameters θ_Y and γ_Y are the elasticities of tax revenue and government spending with respect to output, respectively. These capture what are commonly referred to as "automatic stabilizer" effects. This model is a transformation of the $\eta_t = H\varepsilon_t$ parameterization. The transformations linking the parameters to H are

$$\begin{aligned}
\theta_G &= \frac{H_{12} - H_{32}H_{13}}{1 - H_{23}H_{32}}, \theta_Y = H_{13} \\
\gamma_T &= \frac{H_{21} - H_{23}H_{31}}{1 - H_{31}H_{13}}, \gamma_Y = H_{23} \\
\zeta_T &= \frac{H_{31} - H_{32}H_{21}}{1 - H_{21}H_{12}}, \zeta_G = \frac{H_{32} - H_{31}H_{12}}{1 - H_{21}H_{12}},
\end{aligned} \tag{8}$$

This mapping allows for direct comparison with the TVV-ID results.

4.2 Estimates & tests

In this section, I present evidence in favour of the conditions for TVV-ID to hold, report estimates of the structural parameters based on TVV-ID, and test the identifying assumptions of both BP and MR.

Testing the identification conditions

To motivate TVV-ID, Figure 2 plots "eyeball" evidence of heteroskedasticity in the data using moving averages of squared disturbances for the reduced form residuals, BP's shocks, and MR's shocks in turn; in all three sets of series, there appear to be strong patterns of heteroskedasticity. Table 2 formalizes this evidence using the test proposed in Theorem 3. Specifically, I test the rank of $\widehat{cov\left(\zeta_t, \zeta_{t-1}\right)}$, the estimate of the first autocovariance of $\eta_t \eta_t'$, to assess the identification conditions of Theorem 1. The tests are of the null hypothesis $rank\left(\cdot\right) = r$ against the alternative of $rank\left(\cdot\right) > r$. In this 3-variable system, a rank of 3 implies that $M_{t,t-1}$ satisfies the conditions imposed by Theorem 1, as shown in Proposition

Table 3: Estimates

	BP	MR	TVV-ID
θ_G	-0.06	-0.20	-0.13 (0.10)
θ_Y	2.08	3.13	1.58 (0.18)
γ_T	0	0.06	0.11 (0.13)
γ_Y	0	0	0.02 (0.39)
ζ_T	-0.08	-0.35	-0.00 (0.02)
ζ_G	0.07	0.10	0.06 (0.045)

The first two columns are estimates obtained in Mertens & Ravn (2014). The third column maps estimates of H obtained via TVV-ID to the parameters of BP and MR using (8). The TVV-ID estimates result from fitting the AR(1) SV model, described in Section 3, with details provided in Supplement 2.2.

2. Both tests easily reject ranks smaller than 3 at the 1% level, indicating a rank exceeding 2, so the model is well-identified by TVV-ID.

Estimates of structural parameters

Estimation based on TVV-ID proceeds using the AR(1) SV approach recommended by the simulation study. The estimates are reported in the third column of Table 3, with BP and MR results for comparison.[15] The structural shocks themselves are extremely well-correlated with the BP shocks and very well-correlated with the MR shocks. The one statistically significant parameter estimate is that central to the tax multiplier debate, θ_Y, for which I obtain the value 1.58 with a 95% confidence interval of $[1.23, 1.94]$.

Testing the Blanchard & Perotti (2002) assumptions

The three identifying assumptions made by BP can be directly tested from the estimates of H. First, for the elasticity of tax revenues with respect to output, θ_Y, I obtain a value of 1.58, and can reject BP's calibrated value 2.08 at the 1% level. In the version of their model documented in MR, spending is assumed to respond contemporaneously only to its own shocks: $\gamma_T = \gamma_Y = 0$. In the original paper, $\theta_G = 0$ (taxes do not respond to spending)

[15] It is well-known that EM algorithms can be sensitive to start values; thus, optimization was carried out across a grid of start values and the median estimates were used to initialize a final optimization. The range of estimates across start values is very small, see Table 4 in the Supplement. As an additional check, the estimates from alternative volatility models (same Table) are extremely similar.

is an alternative to $\gamma_T = 0$. None of these exclusion restrictions can be rejected; they are consistent with TVV-ID results.

Testing the validity of Mertens & Ravn's (2014) instruments

MR use the RR shocks as external instruments to identify tax shocks. Like standard instruments, they must be both relevant and exogenous (see Montiel Olea, Stock, & Watson (2016)). Thus, for relevance, I compute first-stage $F-$statistics under both homoskedasticity and heteroskedasticity, and compare them to the corresponding rules of thumb, $F > 10$ (Staiger & Stock (1997)) and $F > 23$ (Montiel Olea & Pflueger (2013)). Under homoskedasticity the value is 4.13 and under heteroskedasticity 1.76; the instrument is only weakly related to the endogenous residual. This suggests there could be a weak identification problem. Table 6 in the Supplement shows that this is true of all alternative narrative measures considered by MR. These results are at odds with the reliability measure they report. That measure of how much variation in the instrument is explained by the structural shock is asymptotically equivalent to the R^2. There are reasons to favour conclusions based on the first-stage $F-$statistic. The reliability measure can only be computed based on estimated structural shocks; instrument validity is *assumed* to obtain these. The $F-$statistic also conveys more information because established thresholds are based on how a deficiency in the first-stage quantitatively impacts bias or size-distortion in the second stage.[16,17]

Using the structural shocks from TVV-ID, I can also test the exogeneity assumption required for the proxy VAR. I test the hypothesis that the coefficients in the regression of the RR shocks on ε_t^G and ε_t^Y are zero. The test rejects at the 5% level for the shocks jointly, driven by a significant negative relationship with ε_t^Y. This suggests that, despite careful construction, the narrative measure has not been fully purged of cyclical behaviour, and still contains endogenous variation in tax revenues. Table 6 in the Supplement repeats the exercise for the alternative shocks in MR; only the series based on the full set of RR shocks (including shocks with implementation lags) does not exhibit endogeneity. The strong negative relationship between the instrument and output shocks implies that, for a tax cut, the estimated impact on output could be biased upwards.

[16] Additionally, MR note that the reliability statistic requires the additive form of measurement error specified in the text. However, it is reasonable to believe measurement error could scale with the size of the tax shock being measured, in keeping with several common forms of heteroskedasticity (in linear regression). The reliability itself also offers no measure of the uncertainty around the relationship between the shocks and instrument. While MR do bootstrap the statistic, it is well-known that bootstrapping procedures may not properly capture variability if weak identification is present.

[17] While the instruments considered here appear weak, Mertens & Montiel Olea (2018) focus on the impact of *marginal* tax rates and construct an alternative instrument based on the RR narrative shocks, scaled based on marginal tax rate changes, which appears to be a strong instrument for the tax rate changes they consider.

Mertens & Montiel Olea (2018) hypothesize that effects of tax changes estimated using instruments are possibly underestimated, arguing that the pro-cyclicality of tax rates and bracket creep potentially push estimates *downwards.* This argument is at odds with the evidence here. However, the negative relationship I estimate can be rationalized. Given the need to parse political motivations behind tax changes in order to classify them in RR, it is possible that in an effort to avoid pro-cyclicality, the time series over-omits ideologically-motivated events that may appear pro-cyclical. Further, given the focus on ideologically-motivated tax changes in the RR shocks, a range of political economy stories linking rising incomes with pro-tax cut governments could explain the negative relationship between the instrument and the output shocks. Finally, a bracket creep story can also work in this direction, with bracket creep a consequence of rising incomes, and tax cuts often coinciding with episodes of bracket creep (Mertens & Montiel Olea (2018) note 1964, 1981, 2001).

4.3 Multipliers

The parameter estimates from TVV-ID lead to important differences in dynamic multipliers compared to previous work. Figure 3 plots the dynamic tax multiplier following the methodology of MR. The shock corresponds to a tax cut of 1% of GDP. 95% confidence intervals are computed using the same wild bootstrap as MR for the reduced form portion of the IRF with the ML variance estimates of the structural parameters, combined via the delta method. The differences compared to the BP and particularly MR results in the first panel are stark. The MR IRF is rejected at all horizons; the BP for horizons up to five quarters. As discussed above, endogeneity of the RR shocks with respect to output shocks could be causing an upward bias for MR. The response on impact is -0.02% (not significant) compared to 0.48% for BP and 1.99% for MR. The peak multiplier occurs later and is lower: 0.86 at eight quarters compared to 1.35% at seven quarters (BP) and 3.19% at four quarters (MR). It suggests a more significant response lag of the economy to tax changes than previous results. The second panel recomputes the IRF for BP using the new elasticity estimated via TVV-ID. The path is virtually identical to the TVV-ID path, mimicking the result when MR do the same using their estimated elasticity. This affirms the finding of Caldara & Kamps (2017) that the elasticity explains virtually all estimated differences in multipliers, and shows that the results of BP can be reconciled with those of TVV-ID via the calibrated parameter, θ_Y.

Figure 4 plots the government spending multipliers. The estimates here are much more similar across approaches, as predicted by more similar values for θ_G and ζ_G. On impact, the multipliers are 0.65% for TVV-ID, just lower than BP (0.69%) and close to MR (0.80%). The maximum response is 0.75% for TVV-ID, compared to 0.81% for BP and 0.96% for MR, all

Figure 3: Response to a tax cut of 1% of GDP

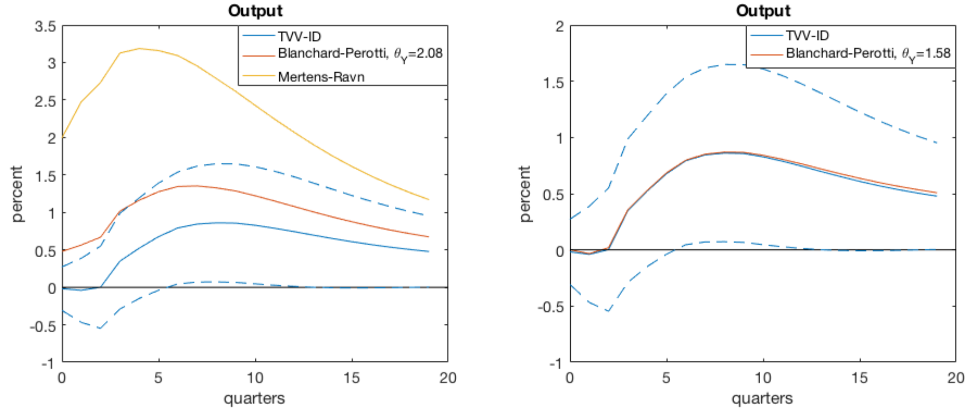

Dashed lines are 95% confidence intervals computed using a wild bootstrap for the reduced form and ML for the structural parameters, combined via the delta method. The BP estimates in the left panel use their elasticity $\theta_Y = 2.08$; the right uses the value of 1.58 estimated via TVV-ID.

Figure 4: Response to a spending shock of 1% of GDP

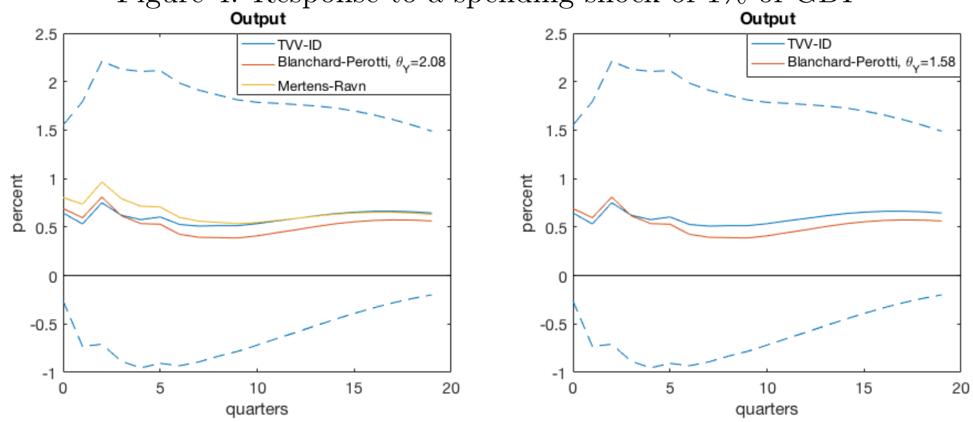

Dashed lines are 95% confidence intervals computed using a wild bootstrap for the reduced form and ML for the structural parameters, combined via the delta method. The BP estimates in the left panel use their elasticity $\theta_Y = 2.08$; the right uses the value of 1.58 estimated via TVV-ID.

after two quarters. The second panel plots the BP response with the elasticity re-calibrated to TVV-ID; doing so barely impacts estimates of the structural parameters, so the paths are virtually identical.

Caldara & Kamps (2017) develop a new methodology using non-fiscal proxies as instruments for identification and find that, in the short run, spending multipliers are larger than tax multipliers. This is also true here, but tax multipliers do eclipse spending multipliers around their peak impact, after two years. Figure 7 in the Supplement plots their estimates against mine. In general, the IRFs are similar, but TVV-ID again yields smaller multipliers on impact. Except for quarters 0-4 for tax cuts, my point estimates lie within their 68% credible sets; all estimates lie within 95% credible sets. They also find that the impact effect largely explains discrepancies in dynamic multipliers across identification approaches, and in turn, that the impacts are governed almost entirely by the elasticities with respect to output and the covariances of u_t (their equation (11)). Accordingly, for their version of BP, using a tax elasticity of 1.7, they obtain an almost identical IRF to mine. For spending, the slightly higher (though not statistically significantly so) spending elasticity I obtain can explain the somewhat lower multipliers obtained. Differences in the shape of the IRFs result from different reduced form coefficients (they use a 5-variable VAR and only a linear time trend). Since the identifying conditions of both schemes hold up to testing (their instruments pass validity tests), it is reassuring that, with the exception of taxes on impact, the dynamic multipliers of each study cannot be rejected under the methodology of the other.

The results under TVV-ID are also in line with the spending multipliers obtained by Ramey (2011b) using a very different methodology based on isolating defense-related spending events. For a sample from 1939-2008, excluding WWII, her estimates range from 0.6-0.8, which includes the values estimated here. Ramey & Zubairy (2018) consider how the multiplier changes across states of the economy using defense spending and BP spending shocks as instruments, which they check carefully for relevance. While in some relevant states of the economy they obtain slightly lower estimates, the 0.6-0.7 range found here accords with their results.

The elasticity of revenues with respect to output

The differences in tax multipliers are largely determined by the lower elasticity of tax revenues with respect to output that I estimate via TVV-ID. This discrepancy between my elasticity and that in the original BP paper may be partially explained by the fact that, in their calibration, they consider data on general government revenue and spending, as opposed to solely federal government. The response of federal revenues should be lower than that of federal, state, and local revenues *combined*. Significantly, Follette & Lutz (2010) develop a more detailed methodology and estimate the elasticity of tax revenues with respect to output for just the federal government, and obtain a value of 1.6 for the period 1986-2008 – nearly identical to what I obtain via TVV-ID. They obtain 1.4 for 1960-1985. On average, their value is thus slightly lower than mine, but they consider a mix of annual data and quarterly data. This accords with BP's argument that lower frequency data will deliver lower elasticities. MR discuss discrepancies between institutional estimates of BP, Giorno, Richardson, Roseveare, & van der Noord (1995), and others and results coming from their instrumental approach; there is far less discrepancy between the results arising from all of the former approaches and my findings based on TVV-ID.

The higher elasticity of MR, 3.13, may result from the weakness of the instrument, as discussed above. In more recent work, Caldara & Kamps' (2017) baseline elasticity estimated using all of their non-fiscal instruments is lower at 2.18. Their instruments pass all tests for relevance and exogeneity, unlike the RR shocks. Their elasticity is still higher than mine, but is based on a different reduced-form VAR. They also show that the high elasticity found in Mountford & Uhlig (2009) – about 3.2 – can be traced to those authors' penalty-function identification approach, which *maximizes* the systematic component of tax revenues.

5 Conclusion

This paper presents a general argument that structural shocks can be identified via time-varying volatility. The previous literature offers identification arguments based on a path of variances available for only a handful of parametric models of the variance process. My identification approach makes minimal assumptions on the variances as a stochastic process. This argument highlights a novel channel of identification based on heteroskedasticity that frees the researcher from needing to assume a particular functional form (or, indeed, any functional form) to obtain identifying moments. This empowers researchers to develop new models and approaches in contexts exhibiting time-varying volatility without needing to re-establish identification for each. Economic information often used to structurally identify such models need only be used to label the shocks identified by TVVID. I propose a simple test of the identification conditions that is valid even when the parameters to be tested are unidentified. A variety of estimation methods are considered. Simulation evidence shows that quasi-likelihood methods based on an auto-regressive log-variance process work well even when the true process has a different form.

My empirical investigation of fiscal multipliers produces estimates that are quite low, but broadly align with previous studies. The tax multipliers estimated by Blanchard & Perotti (2002) can easily be reconciled with TVV-ID by adjusting their calibrated elasticity of revenues. The tax elasticity I obtain, about 1.6, is consistent with the work of Follette & Lutz (2010). For both tax changes and government spending, my results are fairly similar to those recently obtained by Caldara & Kamps (2017). For government spending, my multipliers are consistent with the values in Ramey (2011b) and Ramey & Zubairy (2018). Mertens & Ravn's (2014) high values may result from instrument endogeneity or weakness. These findings contribute to an increasing body of empirical work in favour of multipliers below unity, and to tax multipliers smaller than spending multipliers. This demonstrates the potential of TVV-ID to offer new insights into old problems using an identification approach different to those previously considered in the literature.

A Proofs

A.1 Derivation of Proposition 1

Proof. I start with

$$E_{t,s}\left[\zeta_t \mid \sigma_t, \mathcal{F}_{t-1}\right] = L\left(H \otimes H\right) G\sigma_t^2.$$

Since v_t was shown to be a martingale difference sequence and $\text{Var}_t\left(v_t\right) < \infty$ (Assumption C.2-3),

$$cov_{t,s}\left(v_t, v_s\right) = 0,\ \ s \neq t.$$

This also implies that in the signal-noise decomposition, (5), v_t is white noise. Using this fact, Assumption B, Assumption C.1-2, and the decomposition of ζ_t above, it is immediate that, for $s \neq t$,

$$\begin{aligned} E_{t,s}\left(\zeta_t \zeta_s'\right) = {} & L\left(H \otimes H\right) G E_{t,s}\left[\sigma_t^2 \sigma_s^{2'}\right] G'\left(H \otimes H\right)' L' \\ & + L\left(H \otimes H\right) G E_{t,s}\left[\sigma_t^2 v_s'\right] + E_{t,s}\left[v_t \sigma_s^{2'}\right] G'\left(H \otimes H\right)' L'. \end{aligned} \tag{9}$$

By the law of iterated expectations, Assumption A.1 implies that

$$E_{t,s}\left[\Sigma_t \mid \sigma_s^2\right] = E_{t,s}\left[\varepsilon_t \varepsilon_t' \mid \sigma_s^2\right],\ \ t \geq s.$$

This, in turn, by the law of iterated expectations, implies that

$$E_{t,s}\left[vec\left(\varepsilon_t \varepsilon_t' - \Sigma_t\right) \sigma_s^{2'}\right] = 0,\ \ t \geq s.$$

Thus, setting $t > s$, the third term in (9) vanishes, leaving

$$E_{t,s}\left(\zeta_t \zeta_s'\right) = L\left(H \otimes H\right) G E_{t,s}\left[\sigma_t^2 \sigma_s^{2'}\right] G'\left(H \otimes H\right)' L' + L\left(H \otimes H\right) G E_{t,s}\left[\sigma_t^2 v_s'\right]. \tag{10}$$

Finally, I can rewrite (10) as

$$\begin{aligned} & L\left(H \otimes H\right)\Big(G E_{t,s}\left[\sigma_t^2 \sigma_s^{2'}\right] G' + G E_{t,s}\left[\sigma_t^2 vec\left(\varepsilon_s \varepsilon_s' - \Sigma_s\right)\right]\Big)\left(H \otimes H\right)' L' \\ & \qquad = L\left(H \otimes H\right) G M_{t,s}\left(H \otimes H\right)' L' \end{aligned} \tag{11}$$

where $M_{t,s} = E_{t,s}\left[\sigma_t^2 \sigma_s^{2'}\right] G' + E_{t,s}\left[\sigma_t^2 vec\left(\varepsilon_s \varepsilon_s' - \Sigma_s\right)'\right]$. $M_{t,s}$ is an $n \times n^2$ matrix. Proposition 1 then follows directly. □

A.2 Proof of Theorem 1

I begin by proving two lemmas for properties of the singular value decomposition (SVD).[18]

Definition 1. Define

1. $U_1 D_U U_2' = V$, a reduced SVD, V $n_1 \times n_2$, D_U $d \times d$,[19]
2. C_i is a full rank matrix, $m_i \times n_i, m_i \geq n_i$,
3. $F = C_1 V C_2'$, with $rank\,(F) = d$.

First, I show that a linear relationship exists between the singular vectors of V (U_1, which will later correspond to an unobservable object) and singular vectors of F (which will later correspond to an observable object).

Lemma 1. *There exists a matrix Γ_1 such that $C_1 U_1 \Gamma_1$ is an orthogonal matrix of singular vectors from an SVD of F.*

Proof. Define $Q_1 R_1 = C_1 U_1$, a QR decomposition, and similarly for $U_2 C_2$. Then F can be factored as $F = Q_1 R_1 D_U R_2' Q_2'$. The upper-triangular matrix R_1 is $d \times d$ and full rank since $C_1 U_1$ is full rank d ($rank\,(R_1) \geq rank\,(Q_1 R_1) = rank\,(C_1 U_1) \geq rank\,(F) = d$). Now take another singular value decomposition, this time of $R_1 D_U R_2'$, so $W_1 D_R W_2' = R_1 D_U R_2'$. Then F can be factored as $F = (Q_1 W_1)\, D_R\, (W_2' Q_2')$, which is itself a reduced SVD (it is easily shown D_R are singular values of F, and the corresponding vectors are clearly orthogonal). To obtain Γ_1, recall $Q_1 R_1 = C_1 U_1$ and note $Q_1 R_1 \left(R_1^{-1} W_1\right) = Q_1 W_1$, singular vectors of F, so $\Gamma_1 = R_1^{-1} W_1$, which is guaranteed to exist. □

Definition 2. Define $S_1 D_S S_2' = F$, a reduced SVD.

I now establish the uniqueness of an SVD of F up to orthogonal rotations, accounting for the possibility of repeated eigenvalues.

Lemma 2. *The SVD of F is unique up to rotations characterized by $F = S_1 T_1 D_S T_2 S_2'$ where T_i is orthogonal.*

[18] For a real-valued matrix V, the singular value decomposition $V = U_1 D_U U_2$ decomposes V into two orthogonal matrices U_1, U_2, and a non-negative diagonal matrix D_U. The "singular vectors", columns of U_1 and U_2, are eigenvectors of VV' and $V'V$ respectively. The non-zero singular values (diagonal of D_U) are square-roots of the non-zero eigenvalues of VV' and $V'V$.

[19] A reduced SVD reduces the dimension of U_1, D_U, U_2 to drop singular values equal to zero and their corresponding arbitrary singular vectors.

Proof. For non-repeated singular values in D_S, the corresponding singular vectors are unique up to sign, and the space of vectors corresponding to any k repeated singular values corresponds to linear combinations of any k such vectors. Thus any alternative reduced singular value decomposition of F can be written as $F = (S_1 T_1)\, D_S\, (T_2 S_2')$, since T_i can incorporate any such sign changes or linear combinations. Since $S_i T_i$ must be orthogonal (by definition of an SVD), $T_i' S_i' S_i T_i = I_d$. Then since S_i is orthogonal, $T_i' T_i = I_d$, so T_i is itself orthogonal. □

Definition 3. Define

1. $C_1 = (H \otimes H)\, G$, $n^2 \times n$ with rank n, $C_2 = (H \otimes H)'$, $n^2 \times n^2$ with rank n^2,
2. G is a selection matrix such that $vec\,(ADA') = (A \otimes A)\, Gdiag\,(D)$,
3. $\tilde{S}_1 = C_1 U_1 \Gamma_1 T_1$, singular vectors from any reduced SVD of F,
4. V is $n \times n^2$ and has no proportional rows,
5. $rank\,(V) = d \geq 2$.

Using the relationships I have derived in Lemma 1 between an SVD of the observable F and an SVD of the unobservable V, I now establish conditions under which H is uniquely determined from singular vectors of F. Using Lemma 2, I show that this is true even in the case of repeated singular values.

Proposition 3. *H is uniquely determined from the equations $F = C_1 V C_2'$ provided V has no proportional rows.*

Proof. U_1 is $n \times d$. Note $C_1 U_1 = \left[vec\left(Hdiag\left(U_1^{(1)}\right) H'\right), \ldots, vec\left(Hdiag\left(U_1^{(d)}\right) H'\right)\right]$, where $d \geq 2$ (this follows from the structure of C_1). By the proportional row condition on V, for any j, there exists at least one pair k, l such that $U_{1,j}^{(l)}/U_{1,i}^{(l)} \neq U_{1,j}^{(k)}/U_{1,i}^{(k)}$ for all $i = 1, 2, ..., d,\ i \neq j$. By an argument due to Brunnermeier et al (2017) (the underlying mathematical result also features in Rigobon (2003) and Sentana & Fiorentini (2001)), $H^{(j)}$ is unique up to scale and sign as the right eigenvector of $Hdiag\left(U_1^{(l)}\right) H' \left(Hdiag\left(U_1^{(k)}\right) H'\right)^{-1}$ corresponding to the j^{th} eigenvalue, provided $U_{1,j}^{(l)}/U_{1,i}^{(l)} \neq U_{1,j}^{(k)}/U_{1,i}^{(k)}$. The same argument applies to $C_1 \tilde{U}_1$ where $\tilde{U}_1 = U_1 \Gamma_1 T_1$, provided $\tilde{U}_1$ has no proportional rows. To establish this, take any two rows in U_1 that are not proportional rows; multiplying by full rank Γ_1 cannot decrease their rank (so they do not become proportional). The same holds true for multiplication by the orthogonal T_1. Thus H remains the unique solution to $C_1 \tilde{U}_1$. □

Proposition 4 is re-written in terms of the identifying equations to yield Theorem 1, noting that the requirements imposed on U_1 imply the stated conditions of $M_{t,s}$.

A.3 Proof of Corollary 1

Proof. Corollary 1 follows directly from Proposition 4 above for any column j for which a pair k, l exists such that $U_{1,j}^{(l)}/U_{1,i}^{(l)} \neq U_{1,j}^{(k)}/U_{1,i}^{(k)}$ for all $i = 1, 2, ..., d$. $\square$

A.4 Proof of Theorem 2

Proof. Theorem 2 is based on the argument underlying Proposition 3. Note that the vectorization of $E_t[\zeta_t]$ is given by $vech\left(HE_t[\Sigma_t]H'\right)$, an additional equation of the form found in CU_1. Define $\hat{M} = \begin{bmatrix} U_{1,M} & E_t[\sigma_t^2] \end{bmatrix}$. Then there is an additional column over which to find a k, l pair for j such that $\hat{M}_j^{(l)}/\hat{M}_i^{(l)} \neq \hat{M}_j^{(k)}/\hat{M}_i^{(k)}$ for all $i = 1, 2, ..., d$ $i \neq j$. The no proportional rows condition that applied to $M_{t,s}$ in Theorem 1 now applies to the augmented matrix $\begin{bmatrix} M_{t,s} & E_t[\sigma_t^2] \end{bmatrix}$. Note that this logic can be extended to adding additional autocovariances, etc., in each case adding columns to $\hat{M}$ and thus decreasing the plausibility of the condition failing. $\square$

A.5 Proof of Corollary 2

Proof. Corollary 2 follows directly from Theorem 2 by noting that the $rank\left(\hat{M}_{t,s}\right) \geq 2$ is satisfied given even one dimension of time-varying volatility and considering scenarios under which the proportional row condition fails. $\square$

A.6 Proof of Proposition 2

Proof. I begin by showing $rank\left(cov_{t,s}\left(\zeta_t, \zeta_s\right)\right) = r$ if and only if $rank\left(M_{t,s}\right) = r$. Recall $cov_{t,s}\left(\zeta_t, \zeta_s\right) = L\left(H \otimes H\right) GM_{t,s}\left(H \otimes H\right)' L'$. The elimination matrix L merely deletes repeated rows (and L' columns), so cannot impact rank. Thus it suffices to work with $\left(H \otimes H\right) GM_{t,s}\left(H \otimes H\right)'$. Denote $C = \left(H \otimes H\right)$, which is square with full rank n^2, since H is full rank n. G is a full rank $n^2 \times n$ matrix. First, if $rank\left(M_{t,s}\right) = r$, $rank\left(GM_{t,s}\right) = rank\left(M_{t,s}\right) = r$ since G is rank n. Because C is full rank and square, $rank\left(CGM_{t,s}\right) = rank\left(GM_{t,s}\right) = r$, and likewise $rank\left(CGM_{t,s}C'\right) = r$. Thus, $rank\left(M_{t,s}\right) = r$ implies $rank\left(LCGM_{t,s}C'L'\right) = r$. Going the other way, if $rank\left(CGM_{t,s}C'\right) = r$, then $rank\left(CGM_{t,s}\right) = r$ since C' is full rank and square. For the same reason, it then follows that $rank\left(GM_{t,s}\right) = r$. Because G has rank n, it further follows that $rank\left(M_{t,s}\right) = r$. Thus, $rank\left(LCGM_{t,s}C'L'\right) = r$ implies $rank\left(M_{t,s}\right) = r$, so $rank\left(cov_{t,s}\left(\zeta_t, \zeta_s\right)\right) = r$ if and only if $rank\left(M_{t,s}\right) = r$.

This means that if $rank\left(cov_{t,s}\left(\zeta_t, \zeta_s\right)\right) = r = n$, then $rank\left(M_{t,s}\right) = n$. In that case, $rank\left(M_{t,s}\right) \geq 2$, satisfying the first identification condition. Moreover, since $M_{t,s}$ is $n \times n^2$, it is full rank, so it must have no proportional rows, satisfying the second identification condition. □

A.7 Proof of Theorem 3

Proof. This is a direct restatement of a main result of Cragg & Donald (1996). If $cov\,(\widehat{\zeta_t, \zeta_{t-p}})$ is a consistent and asymptotically normal estimator of $cov\,(\zeta_t, \zeta_{t-p})$, $rank\,(cov\,(\zeta_t, \zeta_{t-p})) < (n^2 + n)/2$ (which it is because the maximum rank of $M_{t,t-p}$ is n) and $cov\,(\widehat{\zeta_t, \zeta_{t-p}})$ is almost surely full rank (which it is due to sampling error in finite samples) then their Assumption 1 is satisfied. Then the Cragg-Donald statistic and its limiting distribution are given in equation (9) of that paper. □

References

Andrews, D. W. K. (1983): "First Order Autoregresssive Processes and Strong Mixing," manuscript.

——— (1993): "Tests for Parameter Instability and Structural Change With Unknown Change Point," *Econometrica*, 61, 821–856.

Angelini, G., E. Bacchiocchi, G. Caggiano, and L. Fanelli (2018): "Uncertainty Across Volatility Regimes," manuscript.

Auerbach, A. J. and Y. Gorodnichenko (2012): "Measuring the Output Responses to Fiscal Policy," *American Economic Journal: Economic Policy*, 4, 1–27.

Barndorff-Nielsen, O. E. and N. Shephard (2002): "Econometric Analysis of Realized Volatility and its Use in Estimating Stochastic Volatility Models," *Journal of the Royal Statistical Society*, 253–280.

Barro, R. J. and G. Liao (2018): "Options-Pricing Formula with Disaster Risk," manuscript.

Bernanke, B. S., J. Boivin, and P. Eliasz (2005): "Measuring the Effects of Monetary Policy: a Factor-Augmented Vector Autoregressive (FAVAR) Approach," *The Quarterly Journal of Economics*, 120, 387–422.

Bertsche, D. and R. Braun (2018): "Identification of Structural Vector Autogregressions by Stochastic Volatility," *Universität Konstanz GSDS Working Papers*, 01.

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

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

Blanchard, O. J. and J. A. Simon (2001): "The Long and Large Decline in U.S. Output Volatility," *Brookings Papers on Economic Activity*, 135–174.

Bloom, N. (2009): "The Impact of Uncertainty Shocks," *Econometrica*, 77, 623–685.

Brunnermeier, M., D. Palia, K. A. Sastry, and C. A. Sims (2017): "Feedbacks: Financial Markets and Economic Activity," manuscript, 1–50.

Caldara, D. and C. Kamps (2017): "The Analytics of SVARs: A Unified Framework to Measure Fiscal Multipliers," *Review of Economic Studies*, 84, 1015–1040.

Campbell, J. Y., S. Giglio, C. Polk, and R. Turley (2018): "An Intertemporal CAPM with Stochastic Volatility," *Journal of Financial Economics*, 128, 207–233.

Carriero, A., T. E. Clark, and M. Marcellino (2018): "Endogenous Uncertainty?" manuscript.

Chan, J. C. C. and A. L. Grant (2016): "On the Observed-Data Deviance Information Criterion for Volatility Modeling," *Journal of Financial Econometrics*, 14, 772–802.

Cogley, T. and T. J. Sargent (2005): "Drifts and Volatilities: Monetary Policies and Outcomes in the post WWII US," *Review of Economic Dynamics*, 8, 262–302.

Cragg, J. G. and S. G. Donald (1996): "On the Asymptotic Properties of LDU-Based Tests of the Rank of a Matrix," *Journal of the American Statistical Association*, 91, 1301–1309.

Darmois, G. (1953): "Analyse Generale des Liaisons Stochastiques," *Revue de l'Institut International de Statistique*, 21, 2–8.

Davidson, J. (1994): *Stochastic Limit Theory: An Introduction for Econometricians*, Oxford: Oxford University Press.

Diebold, F. X. and J. A. Lopez (1995): "Modeling Volatillity Dynamics," in *Macroeconometrics: Developments, Tensions and Prospects*, ed. by K. D. Hoover, New York: Springer, 427–472.

Favero, C. and F. Giavazzi (2012): "Measuring Tax Multipliers: The Narrative Method in Fiscal VARs," *American Economic Journal: Economic Policy*, 4, 69–94.

Follette, G. and B. Lutz (2010): "Fiscal Policy in the United States: Automatic Stabilizers, Discretionary Fiscal Policy Actions, and the Economy," *Federal Reserve Board FEDS Discussion Series*, 43.

Geweke, J. and R. Meese (1981): "Estimating Regression Models of Finite but Unknown Order," *International Economic Review*, 22, 55–70.

Giorno, C., P. Richardson, D. Roseveare, and P. van den Noord (1995): "Estimating Potential Output, Output Gaps and Structural Budget Balances," *OECD Economics Department Working Papers*, 24.

Gouriéroux, C. and A. Monfort (2015): "Revisiting Identification and Estimation in Structural VARMA Models," manuscript.

Gouriéroux, C., A. Monfort, and J.-p. Renne (2017): "Statistical inference for independent component analysis: Application to structural VAR models," *Journal of Econometrics*, 196, 111–126.

Hyvärinen, A., J. Karhunen, and E. Oja (2001): *Independent Component Analysis*, New York: John Wiley & Sons.

Hyvärinen, A., K. Zhang, S. Shimizu, and H. P. O. (2010): "Estimation of a Structural Vector Autoregression Model Using Non-Gaussianity," *Journal of Machine Learning Research*, 11, 1709–1731.

Jurado, K., S. C. Ludvigson, and S. Ng (2015): "Measuring Uncertainty," *American Economic Review*, 105, 1177–1216.

Kilian, L. (2009): "Not All Oil Price Shocks Are Alike : Disentangling Supply Shocks in the Crude Oil Market," *The American Economic Review*, 99, 1053–1069.

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

Kilian, L. and C. Park (2009): "The Impact of Oil Price Shocks on the U.S. Stock Market," *International Economic Review*, 50, 1267–1287.

Kim, S., N. Shephard, and S. Chib (1998): "Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models," *The Review of Economic Studies*, 65, 361–393.

Lanne, M. and H. Lütkepohl (2008): "Identifying Monetary Policy Shocks via Changes in Volatility," *Journal of Money, Credit and Banking*, 40, 1131–1149.

——— (2010): "Structural Vector Autoregressions with Nonnormal Residuals," *Journal of Business and Economic Statistics*, 28, 159–168.

Lanne, M., H. Lütkepohl, and K. Maciejowska (2010): "Structural Vector Autoregressions with Markov Switching," *Journal of Economic Dynamics and Control*, 34, 121–131.

Lanne, M. and P. Saikkonen (2007): "A Multivariate Generalized Orthogonal Factor GARCH Model," *Journal of Business and Economic Statistics*, 25, 61–75.

Leeper, B. E. M., A. W. Richter, and T. B. Walker (2012): "Quantitative Effects of Fiscal Foresight," *American Economic Journal: Economic Policy*, 4, 115–144.

Lindner, A. M. (2009): "Stationarity, Mixing, Distributional Properties and Moments of GARCH(p,q)-Processes," in *Handbook of Financial Time Series*, ed. by T. Mikosch, J.-P. Kreiß, R. A. Davis, and T. G. Andersen, Berlin: Springer, 43–69.

Ludvigson, S. C., S. Ng, and S. Ma (2018): "Uncertainty and Business Cycles: Exogenous Impulse or Endogenous Response?" manuscript.

Lütkepohl, H., M. Meitz, A. Netšunajev, and P. Saikkonen (2018): "Testing Identifiction via Heteroskedasticity in Structural Vector Autoregressive Models," *Discussion Papers of DIW Berlin*, 1764.

Lütkepohl, H. and G. Milunovich (2016): "Testing for Identification in SVAR-GARCH Models," *Journal of Economic Dynamics and Control*, 73, 241–258.

Lütkepohl, H. and A. Netšunajev (2014): "Disentangling Demand and Supply Shocks in the Crude Oil Market: How to Check Sign Restrictions in Structural VARs," *Journal of Applied Econometrics*, 29, 479–496.

——— (2017): "Structural Vector Autoregressions with Smooth Transition in Variances: The Interaction between US Monetary Policy and the Stock Market," *Journal of Economic Dynamics and Control*, 84, 43–57.

Lütkepohl, H. and T. Schlaak (2018): "Choosing Between Different Time-Varying Volatility Models for Structural Vector Autoregressive Analysis," *Oxford Bulletin of Economics and Statistics*, 80, 715–735.

Magnus, J. R. and H. Neudecker (1980): "The Elimination Matrix: Some Lemmas and Applications," *SIAM Journal on Algebraic Discrete Methods*, 1, 422–449.

McCracken, M. and S. Ng (2018): "FRED-MD and FRED-QD: Monthly and Quarterly Databases for Macroeconomic Research," *Federal Reserve Bank of St. Louis Research*.

Mertens, K. and J. L. Montiel Olea (2018): "Marginal Tax Rates and Income: New Time Series Evidence," *The Quarterly Journal of Economics*, 133, 1803–1884.

Mertens, K. and M. O. Ravn (2013): "The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States," *American Economic Review*, 103, 1212–1247.

——— (2014): "A Reconciliation of SVAR and Narrative Estimates of Tax Multipliers," *Journal of Monetary Economics*, 68, S1–S19.

Milunovich, G. and M. Yang (2013): "On Identifying Structural VAR Models via ARCH Effects," *Journal of Time Series Econometrics*, 5, 117–131.

Montiel Olea, J. L., J. H. Stock, and M. W. Watson (2016): "Inference in Structural VARs with External Instruments," manuscript.

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

Nakamura, E. and J. Steinsson (2018): "High Frequency Identification of Monetary Non-Neutrality: The Information Effect," *Quarterly Journal of Economics*, 133, 1283–1330.

Normandin, M. and L. Phaneuf (2004): "Monetary Policy Shocks: Testing Identification Conditions under Time-Varying Conditional Volatility," *Journal of Monetary Economics*, 51, 1217–1243.

Pötscher, B. M. (1991): "Effects of Model Selection on Inference," *Econometric Theory*, 7, 163–185.

Primiceri, G. E. (2005): "Time Varying Structural Vector Autoregressions," *Review of Economic Studies*, 72, 821–852.

Ramey, V. A. (2011a): "Can Government Purchases Stimulate the Economy," *Journal of Economic Literature*, 49, 673–685.

——— (2011b): "Identifying Government Spending Shocks: It's All in the Timing," *Quarterly Journal of Economics*, 126, 1–50.

Ramey, V. A. and S. Zubairy (2018): "Government Spending Multipliers in Good Times and in Bad : Evidence from U.S. Historical Data," *Journal of Political Economy (forthcoming).*

Rigobon, R. (2003): "Identification through Heteroskedasticity," *The Review of Economics and Statistics*, 85, 777–792.

Rigobon, R. and B. Sack (2003): "Measuring the Reaction of Monetary Policy to the Stock Market," *Quarterly Journal of Economics*, 118, 639–669.

——— (2004): "The Impact of Monetary Policy on Asset Prices," *Journal of Monetary Economics*, 51, 1553–1575.

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*, 100, 763–801.

Sentana, E. and G. Fiorentini (2001): "Identification, Estimation and Testing of Conditionally Heteroskedastic Factor Models," *Journal of Econometrics*, 102, 143–164.

Shephard, N. (1996): "Statistical Aspects of ARCH and Stochastic Volatility," in *Time Series Models: in Econometrics, Finance and Other Fields*, ed. by D. Cox, D. Hinkley, and O. Barndorff-Nielsen, New York: Chapman & Hall, 1–68.

Shimizu, S., P. O. Hoyer, A. Hyvärinen, and A. Kerminen (2006): "A Linear Non-Gaussian Acyclic Model for Causal Discovery," *Journal of Machine Learning Research*, 7, 2003–2030.

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

Skitovich, V. (1953): "On a Property of a Normal Distribution," *Doklady Academii nauk SSSR*, 89, 217–219.

Staiger, D. and J. H. Stock (1997): "Instrumental Variables Regression with Weak Instruments," *Econometrica*, 65, 557–586.

Stock, J. H. (2008): "Recent Developments in Structural VAR Modeling," in *NBER Summer Institute: Methods Lectures.*

Stock, J. H. and M. W. Watson (2002): "Macroeconomic Using Forecasting Diffusion Indexes," *Journal of Business and Economic Statistics*, 20, 147–162.

Uhlig, H. (1997): "Bayesian Vector Autoregressions with Stochastic Volatility," *Econometrica*, 65, 59–73.

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