Macroeconomic shocks: A data-driven investigation

Seminar handbook by Gabriel Züllig / University of Copenhagen / Fall 2021[1]

The seminar is devoted to an independent empirical research project with macroeconomic data. Studying the propagation of macroeconomic shocks – for example a monetary policy tightening, increasing government expenditure, or news about technological innovations – is inherently difficult, because the economy is a complex system of endogenous decisions. At the same time, having convincing evidence on causal relationships of macroeconomic variables is crucial when validating structural macroeconomic models or making policy decisions.

1 Goal

To represent dynamic empirical relationships between data series, macroeconomists typically estimate vector autoregressive models (VARs), local projections (LP), or dynamic stochastic general equilibrium (DSGE) models. Your job will be to apply a tool of empirical macroeconomics to study the propagation of a shock of your choosing and implement at least one of the following extensions:

- Identification: Determining the causal effect of x on y is challenging in an economy where everything is interdependent. The identification restrictions discussed are: recursiveness assumption (Cholesky) / sign restrictions (SR) / high-frequency identification (HFI) / DSGE models whose parameters are estimated[2]
- Nonlinearities: Many real-world frictions and constraints imply nonlinear behavior, which means the response of y to x depends on a third variable s. Ways to implement such nonlinearities in empirical models are: Smooth-transition VAR/LP, Interacted VAR.[3]

To get a flavor for why these two things matter, consider the example of a monetary policy shock, which is used as an illustrating example throughout this brief. Theories with nominal rigidities imply that higher interest rates act as demand shocks: A nominal rate hike implies a higher real interest rate and thus a temporarily lower demand by consumers and/or firms. As prices slowly adjust, real output returns back to trend. Figure 1 shows the estimated response of output and inflation to an interest rate hike in (a) a VAR with the most common *identification* scheme (a Cholesky decomposition) contrasted with a more modern approach referred to as high-frequency identification. (HFI follows the same logic as the instrumental variable approach frequently used in microeconometrics. More details on the data used will follow later in this handbook.) Compare this to an impulse response obtained from a linear local projection in (b), which is a popular alternative to estimating the same impulse response function.

Figure 1: Responses of real output (top) and inflation (bottom) to monetary policy shock

(a) VAR / Cholesky and HFI (b) Local projection (c) LP, nonlinear

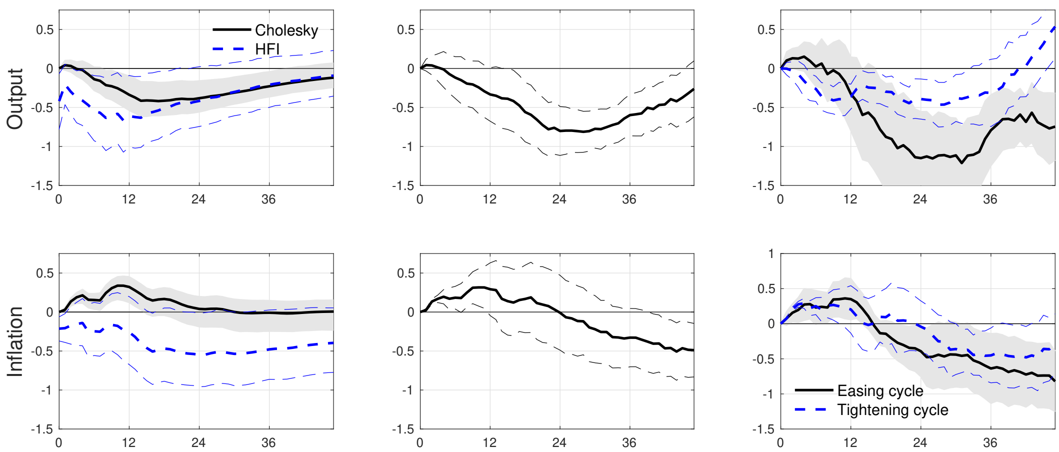

The LP impulse response shows the same behavior over the first few months – this is by construction. In the later periods, however, the contractionary effects on output from the same shock are deeper and more prolonged. A great advantage of local projections is that they are sufficiently flexible to be efficiently extended to a *nonlinear* setting and address a wide spectrum of interesting research questions. For example, panel (c) shows the same impulse responses of GDP and inflation after a surprise monetary policy increase, but this time in two different "states". The blue lines show the response if interest rates have increased in the 12 months prior to the shock, that is if the central bank is in a tightening cycle. In contrast, the black line shows the responses under the assumption that the central bank lowered the interest rate in the 12 months prior to the shock. There is a clear difference in the medium-run response of GDP, whose response is more than twice as strong during an easing cycle. This path-dependence has recently been discussed by Berger et al. (forthcoming) and we can interpret our finding through the lens of their hypothesis: Many mortgages (in the U.S. and many other countries) have fixed interest rates over a 30 year period. When rates decrease, households have the option to re-finance their mortgage and lock in these lower interest rates. When interest rates are on a decreasing path, there is presumably a large fraction of borrowers who can make use of re-financing with a further expansionary shock, freeying up resources for other consumption and investment. As a consequence, demand(GDP) increases substantially. If interest rates are on an upward path, fewer people will decide to re-finance because it would imply larger mortgage payments, and thus the same expansionary shock has smaller effects on GDP.

Clearly, there is still a lot to learn, both in terms of *identification* and *nonlinearities* of effects, even for the most widely studied type of shock in Macroeconomics.

2 Seminar schedule

- 7 September, 10.15-12.00: Kick-off meeting
- Online tutorials
- Digital coffee corners (voluntary): 17 September, 15 October, 29 October at 15.00
- 1 October, 10.00: Submission deadline for project description (1 page; upload to Absalon)
- 5 November, 12.00: Submission deadline for draft (upload to Absalon)
- 15-16 November: Workshop with 20-minute presentations and 15-minute discussions (details follow)
- 1 December, 10.00: Submission deadline for final seminar paper (upload to Digital Exam)

3 Contact

Course instructor: Gabriel Züllig (gz@econ.ku.dk)
All resources are provided via Absalon and my website (gabrielzuellig.com), and these should be your first point of contact. Support meetings generally via zoom, potentially in person, upon request (by e-mail at least 1 day in advance).

4 Shock

When a macroeconomist says "propagation of shocks", she really just means "how does an exogenous shift in x affect y over time?" Therefore, we have to pick an x and a y. Some widely studied x's are:

- Monetary policy shock: discussed at length to illustrate the models introduced throughout this handbook. A good overview is provided by Ramey (2016). If you prefer a history of economic thought on monetary policy shocks and methodological advances in video format, I recommend Jón Steinsson's public lecture at the 2020 American Economic Association Annual Meeting, available here and here. There are interesting empirical studies on "unconventional" policies such as quantitative easing, too (Dedola et al., 2021, Lhuissier and Nguyen, 2021).

- Fiscal policy shock: Fiscal policy tends to significantly increase real output (while crowding in private consumption, such that the fiscal multiplier is > 1 (Blanchard and Perotti, 2002, Mertens and Ravn, 2014)). A more recent instrument based on "military news" questions this result (Ramey and Zubairy, 2018), and the debate on whether fiscal multipliers are higher in recessions (and thus a countercyclical marginal propensity to consume) is still a vivid one (Auerbach and Gorodnichenko, 2012a, Caggiano et al., 2015, Ramey and Zubairy, 2018).
- Financial shock: Gilchrist and Zakrajšek (2012) show that innovations in financial variables that are orthogonal to current economic conditions do dampen real variables in the future, and several similar studies on macro-financial linkages were published in the aftermath of the global financial crisis. What is still a somewhat unsettled issue is the effect on prices, i.e. whether unexpected changes in financial conditions look more like demand or supply shocks (Abbate et al., 2020).
- Uncertainty shock: If economic outcomes become more uncertain, businesses postpone investment and hiring, resulting in lower demand (Leduc and Liu, 2016). Recent evidence points to a great deal of state dependence in this demand channel (Caggiano et al., 2015).
- Shocks to TFP or future TFP ("news shock") TFP disturbances are typically examined using long-run restrictions, where the shock to be identified is the only one that can affect output or labor productivity in the long run (Galí, 1999). This identification is relatively straightforward to implement but not discussed in this handbook. More common nowadays is the identification of news about future technology: They drive stock prices contemporaneously, but real output only with a lag. Beaudry and Portier (2006) find that these shocks explain a large part of macroeconomic fluctuations. An overview of the subsequent literature is provided by Ramey (2016).
- Other potential shocks, including a recent reference as a starting point: exchange rate (Caselli and Roitman, 2019), oil price/cost-push (Känzig, 2021), immigration (Furlanetto and Robstad, 2019), ...

The shock you pick should, first and foremost, catch your interest. If you care about financial markets, pension funds, or public debt, pick a shock that matters for them. If you care about macroeconomic policy (because you want to work for the Ministry of Finance, for example), pick one of the first two. If you're into International Econ, either or both your x or your y could be related to the trade literature (relative prices, trade flows, or interest rate differentials).

Second, the choice should be motivated by economic theory: Why is it interesting, what can we learn from it, or how can you contribute to a better understanding of macroeconomic dynamics?

Keep in mind that these models are typically more suitable to study short-run fluctuations. For questions like "What drives economic growth?", other methods are preferable. Models like VARs represent an inherently stable system (like a steady state).

5 Identification

The next step is to establish the causal relationship from x to y, i.e., to make sure that we get an empirical estimate that is not driven by y's potential impact on x or their joint determination by another factor. Let's start by putting y and potentially other related macro variables (confounding factors) into a vector Y and describe its behavior over time as follows.

$$Y_t = \sum_{j=1}^{p} A_j Y_{t-j} + u_t \qquad E(u_t) = 0,\, E(u_t u_t') = \Omega \tag{1}$$

Y has T (time periods) rows and n (number of variables) columns, and the regressors are the same variables lagged. Typically, we also include a vector of constants, or even a time trend, but they are omitted here. A_j is a $n \times n$ matrix that describes how each row-variable reacts to the j'th lag of the column-variable. Since we have p lags, we have p such matrices. It is common to write A as a function in the lag operator:

$A(L) \equiv I - A_1 L - A_2 L^2 - ... - A_p L^p$, because this allow to write the VAR as follows:

$$\begin{aligned} Y_t &= A_1 Y_{t-1} + ... + A_p Y_{t-p} + u_t \\ Y_t - A_1 Y_{t-1} - ... - A_p Y_{t-p} &= u_t \\ A(L) Y_t &= u_t. \end{aligned} \tag{2}$$

u is a matrix of $T \times n$ reduced-form residuals, i.e., innovations in Y that are not explained by the lags of all the variables in Y. Because we estimate the model with OLS, their expected value is zero and their variances and covariances are described by the matrix Ω.

So far, this model is not identified, and we are simply describing correlations. After estimating Equation 1, we would for example know that in quarter t, the interest rate was 0.1 higher than what historical data would have us expect, but we don't know the nature of the underlying structural shock: Was it because the central bank decided to surprise the economy with a tightening (monetary policy shock), or for example because there was an oil price shock which pushed up prices and the central bank increased the interest rate as they usually do (Taylor rule). The remaining subsections in 5 are devoted to establishing this causality.

5.1 Recursive identification / Cholesky decomposition

The problem of an unidentified VAR is that the u's are correlated. For example, every "shock" to the interest rate implies a correlated innovation in all the other variables, rendering identification impossible. Our goal is to find a matrix S which, when multiplied with u, gives us a vector of uncorrelated errors, giving us a structural interpretation. Mathematically, the transformation and necessary condition for identification are

$$u_t = S\varepsilon_t \tag{3}$$

$$E(\varepsilon_t \varepsilon_t') = I \tag{4}$$

In words: Every reduced-form residual is a linear combination of the structural shocks ε. And these structural shocks are uncorrelated, i.e. the off-diagonal elements of its variance-covariance matrix are zero. Thus, the goal is to find an S that satisfies both conditions. How do we find that S? We start from the variance-covariance matrix of the u's, which is Ω, and use the above mapping to write

$$\begin{aligned} E(u_t u_t') &= \Omega \\ E((S\varepsilon_t)(\varepsilon_t' S')) &= \Omega \\ S \underbrace{E(\varepsilon_t \varepsilon_t')}_{\equiv I} S' &= \Omega \end{aligned} \tag{5}$$

Therefore, we are left with the condition $SS' = \Omega$, and we can recover the matrix S by means of a Cholesky decomposition (i.e., we look for a lower-triangular matrix which, multiplied with its transpose, is equal to the variance-covariance matrix of the reduced-form VAR). We can do so in a single line of computer code: $S = \text{chol}(\Omega)$. Let us keep in mind that S is a lower-diagonal matrix, i.e., has zero-entries 'northeast' of the diagonal. (This has important implications which will become clear shortly.) Once we have S, we can directly retrieve the vector of historical shocks: $\varepsilon\,(= S^{-1}u)$, i.e. we will have an estimate of a monetary policy shock at each point in time.

Impulse response function Perhaps more interesting from the point of view of a researcher interested with the general dynamics of the macroeconomy, we can compute impulse response functions, i.e., estimate how all the variables in Y typically react to a certain shock. Take for example the x^{th} column of the ε_t matrix, so the impulse response function tries to get at $\frac{\partial Y_t}{\partial \varepsilon_t^x}$. To get there, we make use of the "inverted" VAR(p) process describing the variables in Y as a function of previous shocks in Equation 2. For simplicity, I will illustrate the process assuming that we only have one lag $p = 1$. Remember that L is the lag operator, so $LY_t = Y_{t-1}$.

$$\begin{aligned} A(L)Y_t &= (I - A_1 L)Y_t = u_t = S\varepsilon_t \\ Y_t &= (I - A_1 L)^{-1} S\varepsilon_t \quad \text{or more generally} \\ &= \underbrace{A(L)^{-1}}_{\equiv C(L)} S\varepsilon_t = C(L)S\varepsilon_t \end{aligned} \tag{6}$$

In other words, as soon as we have the reduced-form coefficients of the A_j's, we can calculate the reduced-form impulse responses C,[4] multiply by the Cholesky factor S, and obtain the structural IRF we are interested in ($C(L)S$). One crucial implication: Because S has zero entries above the diagonal, some of these responses will be zero. More concretely, the shock in position x has by construction zero impact on all variables ordered before (from position 1 to $x - 1$). These zero restrictions are econometrically necessary in order to achieve identification, but they can be economically problematic because you have to assume that for example GDP does not react to monetary policy decisions within a quarter.

Monetary policy shock The literature on monetary policy shock typically follows the following rule: "real" or "slow" variables like output or unemployment are ordered first. The nominal interest rate is ordered as late as possible because it is "the most exogenous" one. Fast-moving financial variables that are likely to respond immediately such as stock market prices or risk premia, if included, come last. Let's set up the following monetary VAR with 3 variables (GDP, inflation, and the nominal interest rate), a constant, and 1 lag.

$$\begin{pmatrix} y_t \\ \pi_t \\ i_t \end{pmatrix} = \begin{pmatrix} c^y \\ c^\pi \\ c_i \end{pmatrix} + \begin{pmatrix} a^{yy} & a^{y\pi} & a^{yi} \\ a^{\pi y} & a^{\pi\pi} & a^{\pi i} \\ a^{iy} & a^{i\pi} & a^{ii} \end{pmatrix} \times \begin{pmatrix} y_{t-1} \\ \pi_{t-1} \\ i_{t-1} \end{pmatrix} + \begin{pmatrix} s^{y\varepsilon^y} & 0 & 0 \\ s^{\pi\varepsilon^y} & s^{\pi\varepsilon^\pi} & 0 \\ s^{i\varepsilon^y} & s^{i\varepsilon^\pi} & s^{i\varepsilon^i} \end{pmatrix} \times \begin{pmatrix} \varepsilon_t^y \\ \varepsilon_t^\pi \\ \varepsilon_t^i \end{pmatrix} \tag{7}$$

The first row of Equation 7 says that GDP (or some deviation of it from a steady state) responds with a coefficient a^{yy} to its own lag, $a^{y\pi}$ to the lag of inflation, and so on. It also stipulates that the reduced-form residual of the GDP series is a linear combination of ε_t^y – which we could for example call a demand shock – and zero. Compare this to the third row: The last row of matrix A describes the monetary policy reaction function, i.e., how i usually responds to output and inflation. This is what central bankers refer to as the Taylor rule. On top of that, interest rates can deviate from this due to all three structural shocks, among which of them is the monetary policy shock ε_t^i. It also becomes clear that y_t and π_t cannot contemporaneously react to the monetary policy shock. This is what identifies the model, but also where the "recursiveness" of this approach comes from. Depending on the application, this can be problematic, but for monetary policy, it is the standard way to go. Crucially, the restriction only holds for the period of the shock itself. If there is a positive value in ε_t^i, it will directly change i_t (not y_t or π_t), but changing that i_t means that, through the dynamic matrix A, we will change y_{t+1}, a.s.f.

The three time series that enter the model are the following: For "output", a monthly series is not as straightforward to obtain. First, I construct a monthly series that has the properties of (quarterly) real GDP based on two higher-frequency estimates of the Chicago Federal Reserve and IHS Markit, a private data provider. I then compute the percentage deviation from a trend estimate ("potential GDP") provided by the US Congressional Budget Office. Therefore, the series can be interpreted as an output gap measure in percentage points. (The online tutorial on data compilation and organization contains more information.) The second series ("inflation") is the 12-month change in the Core CPI, a price index that includes a typical consumption basket excluding the volatile energy and food prices. The baseline interest rate is the monthly average of the Federal funds rate (the money market rate targeted by monetary policy of the Federal Reserve System). The full $C(L)S$ for a horizon of up to 4 years are displayed in Figure 2.

These impulse responses can be interpreted as follows: If the central bank decides to implement a monetary policy that unexpectedly raises the policy rate by 1 percentage point, real activity falls by 0.4% within 1 or 2

Figure 2: Cholesky

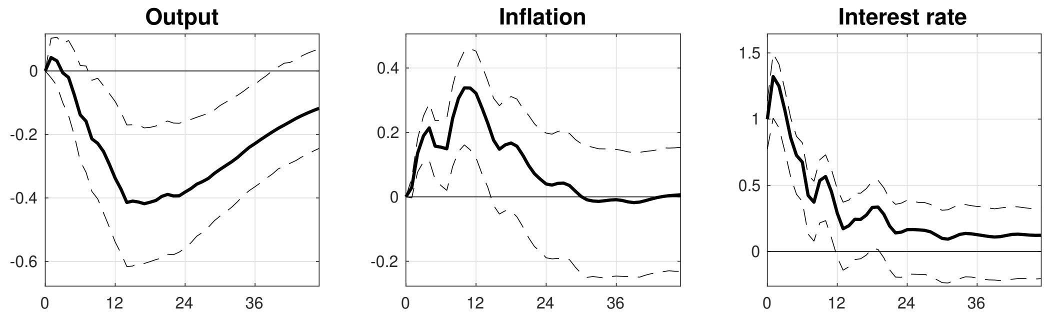

years and reverts back to trend within 4 to 5 years. This is in line with what a New Keynesian model would predict: Because prices are sticky, the real interest rate increases with the nominal one, after which households postpone consumption and economic activity slows down. At the same time, however, prices do not respond the way the model would predict: They increase, even significantly, in the short run. This is common to small VARs identified with the recursiveness assumption and referred to as the "price puzzle". A great deal of work has been devoted to resolving this puzzle. For example, it has been shown that controlling for commodity prices and more monetary variables can alleviate (if not entirely solve) the problem (Christiano et al., 1999). More generally, it could be that Y and π alone do not reflect the full set of information when central banks make decisions on their policy stance. In particular, the price puzzle could be explained by the fact that the Fed has more information on future inflation at the time it increases the policy rate (Sims, 1992, Bernanke and Boivin, 2003). Ramey (2016) provides an overview of the history of thought on monetary policy shocks since, some of which is reflected in this handbook.

Furthermore, one should keep in mind that the choice of the ordering has to be justified. Here, we assume that neither prices nor output can react to monetary policy shocks on impact because they are rather slow-moving variables. An alternative ordering would be less convincing: If we were to order i first, that would mean that monetary policy never reacts to developments in the economy in the same period, which is at odds with what we know of how central banks operate: they constantly monitor a large amount of economic data in real time. If you apply a Cholesky decomposition in your own project, you have to justify the sense of these restrictions economically, even if it is the default in the literature. Because it is easy to implement – on literally one line of code – it is always good that if you choose a different identification strategy as your baseline, to compare it to one with a Cholesky decomposition.

Confidence intervals A few words on how to obtain the "standard errors" used to construct the confidence intervals depicted in Figure 2. Since our biggest concern is that we are not accurately estimating the errors and therefore S, we do the following: From our initially estimated model, we take the matrix A and the empirical distribution of u. We draw from this distribution and generate our own "artificial" data by simulating forward using A. Based on this data, we re-estimate the model and the impulse responses. We do so 1000 times and then obtain the confidence interval, i.e., the 5^{th} and 95^{th} percentile, which are depicted using the dashed lines.

Pros:

- Most established / widely used ("default")
- Computationally easy in many models and more robust than alternatives
- Results are often a good first approximation of what more elaborate identification strategies will show.
- Can potentially identify multiple shocks at once, effortlessly recover the underlying series ε and/or compute forecast error variance decompositions.

Cons:

- Very limiting in less known shocks; The "story" to justify recursive identification in a particular case often has to do with timing, e.g. when the government increases its spending G, it cannot have reacted to the latest GDP figures because it takes time to implement fiscal policy.
- For the most widely studied shock (to i), it generates a short-run response of π that is counterfactual to most structural models ("price puzzle").

5.2 Sign restrictions

An alternative to choosing S relies on making certain restrictions to $C(L)S$ in Equation 6 rather than pinning its contemporaneous value to exactly zero. Suppose that at any point in time, the economy is hit by a linear combination of shocks, all of which have uniquely different implications for the *sign* of the response of all the variables. As an illustrating example, consider that the economy experiences 3 shocks, based on which we expect the 3 baseline variables to expect as follows:

Table 1: Identifying assumption for impulse response functions

		Structural shock		
		Monetary pol. (ε^i)	Demand (ε^y)	Cost push (ε^π)
Effect on	y_t	$-$	$+$	$-$
	π_t	$-$	$+$	$+$
	i_t	$+$	$+$	$+$

How does this identify the monetary policy shock? We can say that any moment in which output and inflation move in the same direction, it cannot be supply shock (which has the opposite implication). On top of that, we distinguish between a monetary policy and other demand shocks (e.g. government spending) by restricting movements of the interest rate. We call equally-sized movements in output and inflation that is accompanied by opposite-signed reaction of the interest rate a monetary policy shock. These predictions are in line with the New Keynesian model and a Taylor rule: If inflation increases after a demand shock, the central bank raises the interest rate. But when the central bank exogenously raises the interest rate, inflation goes down.

Using this rationale, we can disentangle reduced-form residuals and assign each u a weight of the three shocks (ε^y, ε^π, ε^i) based on how the variables evolved thereafter. Mathematically, we are looking for a matrix Q which is orthonormal: multiplied by its transpose, it returns the identity matrix (which is the necessary condition for identification).

$$E(\varepsilon_t \varepsilon_t') = I = Q'Q = QQ'$$

If Q describes the variances-covariance matrix of ε, those structural shocks are identified. At the same time, we recall from equation 5 that the Ω matrix is decomposed into a *lower triangular* matrix S multiplied with its transpose (the Cholesky decomposition). Let us, as a conjecture, define $P = SQ$, so we can show that

$$PP' = SQ(SQ)' = S \underbrace{QQ'}_{\equiv I} S' = SS' = \Omega$$

This shows that multiplying the lower-diagonal S with an orthonormal Q results in a P, which still has the property of describing the variances and covariances of the u's. Therefore, we have shown that P is a plausible set of weights mapping $u_t = P\varepsilon_t$ without the tight "corset" of some entries being restricted to zero.

Computationally, we need to do the following:

1. Find S: Compute the lower-diagonal Cholesky matrix of Ω (not new).
2. Find a Q: There are many possible orthonormal matrices. To obtain one, draw a matrix of random variables from the $N(0, I)$ distribution and perform a so-called QR-decomposition. This will result in two factors: an orthonormal matrix candidate $Q^{(i)}$ and an upper-diagonal matrix $R^{(i)}$ (which we do not need).
3. Calculate $P^{(i)} = SQ^{(i)}$ and compute the structural impulse response functions using the orthogonalization $(C(L)SQ^{(i)})$.
4. Test if the resulting IRFs satisfy the restrictions outlined in Table 1. If not, discard the $Q^{(i)}$ candidate matrix and draw a new one. If all restrictions (or their exact inverse) are satisfied, keep the candidate.
5. Repeat steps 2 to 4 a large number of times, recording all Q's that satisfy the restrictions. They comprise the set of admissible structural VAR models.

Monetary policy shock Figure 3 displays the impulse responses when the dynamic structural restrictions are implemented as proposed in Table 1. The restrictions are set to hold for 4 quarters after the shock, but this setting is relatively flexible. We can for example say a particular restriction has to be satisfied in month 20 after the shock. In the resulting impulse response, the effect on inflation is now negative. We got rid of the price puzzle, but that is at least for the short run by constructions (because we assumed it). This is one of the biggest drawbacks of sign restrictions: Many things have to be assumed, and its often difficult to find a smart combination of restrictions that a.) identify the shock, i.e. pin down the necessary variables to disentangle the sign of responses but b.) keep at least one variable of interest unrestricted so we can actually say something about it. What Figure 1 *does* allow us to say is that the negative inflation response is much more persistent, particularly in light of an output response that is barely significant. This is a common feature to this type of identification (Canova and Nicoló, 2002, Uhlig, 2005).[5] Second, error bands are considerably wider. We have a considerable degree of "shock uncertainty" (shown by wider bands even on the interest rate), but it could also point to the fact that we still have issues identifying our shock properly (Fry and Pagan, 2011).

Figure 3: Sign restrictions

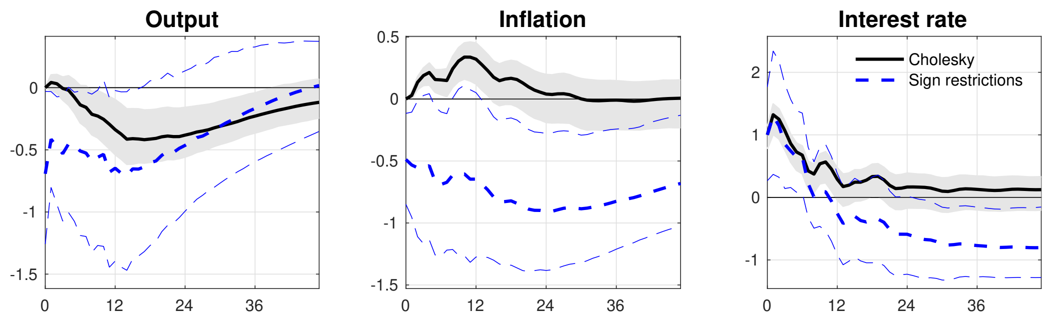

Sign restrictions have fallen a bit out of fashion, although interesting innovations are still being made. You can, for example, impose restrictions not on the IRF but on any other property of the model. If you know that there was a monetary policy shock on a certain date – because the governor gave a widely cited speech – you can require the value of ε_t^i to be negative at a particular t. Or you can restrict the ratio of two variables to be above a certain threshold at time $t + h$. See Antolin-Diaz and Rubio-Ramirez (2018) or Abbate et al. (2020) for some inspiration.

Good sign restrictions are often rooted in economic theory. A simple example: We know that the output price of a firm is equal to a markup over marginal cost, but since neither of the two are observed in reality, whenever there is inflation we don't know which of the two is responsible. Cost-push (think oil price hike) and markup shocks (think increase market power due to competitor exit) have the same implication for the direction of inflation (and real output). The former implies that the firm's profit decreases, but in the latter they should increase. We could use a measure of firm profits, stock market valuation or the like to disentangle the two shocks.

Pros:

- No responses are restricted to be exactly zero.
- Justifications for restrictions tend to be more rooted in economic theory than the ordering in Cholesky-identified SVAR.
- Not necessary to bootstrap on artificial data because we can implement confidence intervals directly from the uncertainty of the Q draws.

Cons:

- Often difficult to find identifying restrictions, meaning they pin down one and only one possible shock ("partial identification").
- Results can be a direct product of the researcher's own assumptions.
- Wide confidence bands (direct result of shock uncertainty)

5.3 High-frequency identification / Instrumental variable approach / Proxy-VAR

In recent years, macroeconometricians have rifled through the toolbox of their micro colleagues and thought about how to apply the instrumental variable approach in a time series context. The idea of an instrument z is that it exogenously shifts the variable x but is uncorrelated to y (other than through x), such that we can identify the causal impact of x on y. As in other fields, finding suitable instruments is not always a straightforward endeavor. In the context of monetary policy shocks, researchers typically rely on short-term movements of financial variables around certain events. They look at movements of rates or yields during relatively narrow windows around monetary policy announcements. The assumption is that prior to the announcement, the market has priced in expectations of how the policy rate should move, given the state of the economy. If yields move up during the window, it represents a monetary policy more contractionary (or less expansionary) than anticipated, and this surprise can be used within a VAR framework.

Let's start from the familiar definition that our reduced-form residuals are a linear combination of the structural, orthogonal shocks ε. Since the IV approach can only identify one shock at a time (contrary to Cholesky or sign restrictions), we write what was a matrix S before as a vector s of dimensions $n \times 1$. We can think of it as the column of the S matrix that corresponds to the shock's same-period impact on each of the variables represented in the VAR.

$$u_t = s'\varepsilon_t^i$$

Further, let z_t be an instrumental variable (or several) which is orthogonal to all potential other shocks in the economy, denoted ε_t^j for $j \neq i$. To be a valid instrument, z must be correlated with our shock of interest but not with others. The exclusion restriction is relatively straightforward in the case of HFI because it is unlikely that other shocks have any systematic effects on the financial variables during those narrow time windows.[6]

$$E[z_t\varepsilon_t^i] \neq 0$$
$$E[z_t\varepsilon_t^j] = 0$$

We're going to proceed similar to two-stage least squares procedure in the cross-section: In the first stage, we regress the potentially endogenous reduced-form residual associated with monetary policy, i.e. u_t^i on the instrument and generate the predicted values

$$u_t^i = \beta_1 z_t + \zeta_t \tag{8}$$
$$\hat{u}_t^i = \hat{\beta}_1 z_t$$

In the second stage, we regress the other variables j on the predicted values in $\hat{u}_t^i$, which now only contains exogenous information on ε_t^i, i.e. is orthogonal to ξ_t and gives a consistent estimate of β_2

$$u_t^j = \beta_2 \hat{u}_t^i + \xi_t = \frac{s^j}{s^i}\hat{u}_t^i + \xi_t \tag{9}$$

Crucially, any contemporaneous movement in Y_t^j is due to u_t^j, and according to 9 any changes in those residuals due to the structural shock of interest is contained in $\hat{u}_t^i$. The respective second-stage coefficients therefore identify the within-period effects of Y^j with respect to ε^i, denoted s^j, up to the size of s^i.[7] Impulse response functions for higher horizons are found by simulating the system further using the familiar dynamic multipliers (equations 1 and 6).[8]

Monetary policy shock The application of a Proxy-VAR to the monetary policy environment is a close replication of Gertler and Karadi (2015): Apart from the previously used variables, we need an instrument z, which is in this case the change of the three-months fed funds future during a 30-minute window around FOMC announcements. Often, such an instrument series will not be available for the entire time period of the VAR sample. Because the IV runs in two separate stages, one simply selects the available subsample to identify the shock (Equation 8). The second step (Equation 9) and simulation of IRFs for later horizons use the u's and the A matrix estimated on the full sample. A second new ingredient is a measure of corporate risk premia, which is a US bond spread stripped from its "fundamentals" (the default probability) and thus represents a notion of investor sentiment or risk appetite. An increase in the excess bond premium means that given the same probability of default, investors demand higher compensation for risk. The indicator is developed in Gilchrist and Zakrajšek (2012). It is included to show that Proxy-VARs are particularly suited to include financial variables (as opposed to Cholesky decompositions, where the lagged responses of financial variables are difficult to justify). Second, one wants to be sure that our u^i are not contaminated with information contained in fast-moving financial variables.

Figure 4: Instrumental variables: High-frequency identification

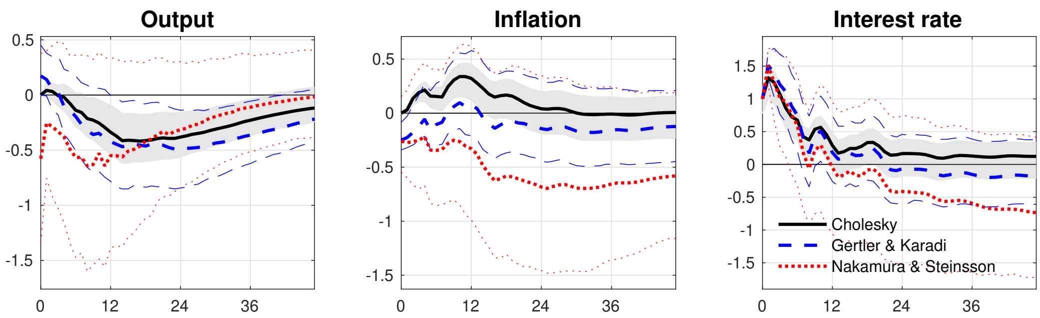

The results in Figure 4 show that the "price puzzle" has vanished. Prices actually respond negatively in the very short and more so in the medium run, although the responses are still muted, while the evidence points to the fact that output effects are at least as strong as when identified recursively. When repeating the exercise with a similar series incorporating information at many more points along the yield curve and provided by Nakamura and Steinsson (2018), results become again more blurred. This is likely due to the shorter time availability of their instrument, rendering identification less powerful, or due to the "central bank information effect". The reader is referred to the footnotes for a starter on this very recent literature.

Confidence intervals The standard bootstrapping algorithm includes creating "artificial" data based on draws from the empirical distribution of u. This changes the t-ordering of residuals, which we should refrain from in a Proxy-VAR because we cannot simply reorganize the timing of z. Instead, one relies on a version of "Wild bootstrapping": Multiply the estimated errors in their original sequence by a random variable drawn from the Rademacher distribution, i.e., randomly flip the sign of all the u's and time-corresponding z with a probability of one half. Construct the counterfactual data and re-estimate. Do this many times over and define the 5^{th} and 95^{th} percentile for the lower and upper bound of the confidence bands.

Narrative approach There is a simpler, yet less clean way to get hold of a z which is used further below and follows the following idea: Our goal is to purge the intended endogenous response of interest rates to the state of the economy – a (version of the) Taylor rule – from changes in the policy rate. Romer and Romer (2004) get this by regressing the change in the Fed's target rate on information about economic fundamentals and expectations about the future at the time of the decision. The residual of the following regression is then defined as the monetary policy shock, which we could use as the z in the Proxy-VAR.[9]

$$\Delta i_t = \beta_1 i_t + \beta_2 u_t + \beta_X X_t' + \varepsilon_t^i \tag{10}$$

Let the vector of controls X be comprised of the following variables from the "Greenbook", the Federal Reserve's own account of expectations: the GDP growth rate and inflation rate, both over the past, current, and subsequent two periods, as well as the revisions of all these forecasts. The resulting series of shocks has the appealing feature that it is "surprising" to the central bank and – by construction – orthogonal to the bank's assessment at the time of implementation. Because they tell a relatively simple story of a true "shock", they are often referred to as "narrative shocks". Unfortunately, the result itself is less convincing: Despite using the narrative shocks as an instrument, the impulse responses look similar to a simple Cholesky decomposition.

Figure 5: Instrumental variables: Narrative shocks

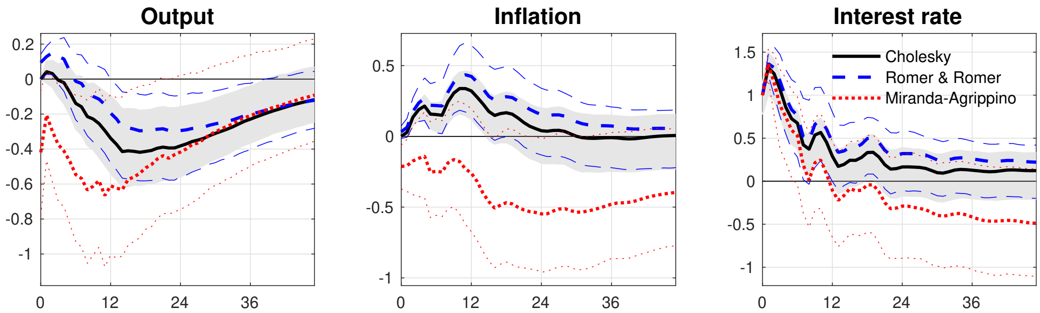

However, Miranda-Agrippino (2016) has combined the best of both worlds: Taking the high-frequency interest rate movements from Gertler and Karadi (2015) and filtering the information effect by applying 10 not only makes output respond fast and stronger, it also renders the monetary policy shock unequivocally deflationary at least in the medium run.

Finding a good instrument As in the cross-sectional application of the IV approach to identification, finding a good instrument is not always as easy as it seems. We are looking for a *series* of information that shift x but has no direct effect on y. We don't necessarily need the information on *each* point in time, but at least on several and fill the remaining time periods with zeros. For example, Piffer and Podstawski (2017) use the change in the gold price around a few global events related to uncertainty – the most prominent being the 9/11 terrorist attacks – to help identify uncertainty shocks. Natural disasters or weather occurrences are other events where exogeneity is easy to justify. High-frequency information from financial markets or Google search statistics can be helpful, easily accessible sources. Känzig (2021) uses scraped press releases by the OPEC to filter out information that drive the oil price. A different line of argument could be to use "exogenous" information on a foreign country. Fiscal (Mertens and Ravn, 2014) or monetary events (discussed above) in the U.S. could affect the bilateral exchange rate between the USD and the currency of a country of interest in a way that is orthogonal to what's going on in that country (although keep in mind the exclusion restriction: the instrument then should affect the domestic economy *only* through the exchange rate). If you run a Proxy-VAR, try to be creative – this seminar explicitly encourages you to take risk and explore original ideas, even if things don't work out in the end.

Pros:

- Can best distinguish between true monetary policy surprises and interest rate movements that are endogenous (e.g. Taylor rule) and therefore anticipated.
- Few lines of code and therefore easily applicable to many models and contexts.

Cons:

- Appropriate instruments can be hard to find, and time series are often shorter.
- u's have to be well-specified to begin with. The procedure does not forgive models that are not well-specified.

This concludes the discussion of the most common ways to identify shocks in empirical models – zero restrictions on impact, sign restrictions, and external instruments/narrative strategies. It should be noted that it is not uncommon to combine several strategies in one model.

6 Models

6.1 Linear local projection

Local projections as proposed originally by Jordà (2005) are a generalized, but also computationally simplified way of estimating impulse response functions. Rather than deriving the dynamic multipliers as in Equation 6, we estimate them directly:

$$Y_{t+h} = \Gamma_h Y_t + \sum_{j=1}^{p} A_j Y_{t-j} + u_{t+h} \tag{11}$$

Γ_h is an $n \times n$ matrix describing how each variable loads on its own and lags of other variables between point t and some point in the future $t + h$. If the output measure is ordered first and the "shocked" interest rate third, then the coefficient of interest will be the first-row, third-column element of Γ_h, defined γ_h. Performing this regression for different horizons h between 0 and H gives us a vector of length $H + 1$ with many $\hat{\gamma}_h$'s. Stacked together, they describe the impulse response function. Plagborg-Møller and Wolf (2021) show that for $h \leq p$, this impulse response is the same as the one from a VAR. For larger horizons, the IRFs are typically less persistent and less "stable" because of the fact that the model imposes no dynamic structure whatsoever. You can think of this as both a drawback (because impulse responses look less "nice") or an advantage (because the VARs impose a lot of structure into how economic variables dynamically relate to one another, and local projections are more "honest" representation of these dynamics).

Note that local projections are merely a more direct way to estimate the impulse response function (denoted C in Section 5.1), but the challenges in terms of identifying the shock itself still apply. Like in a VAR, we could again rely on recursiveness assumptions to do so or use the instrumental variable approach.

In the first case, we estimate the contemporaneous impact matrix S as usual as the orthogonalized variance-covariance matrix from the reduced-form VAR residuals. The second case is when the shock can be identified externally as with high-frequency identification. Then we can estimate γ_h even more directly by

$$Y_{t+h} = \gamma_h z_t + \sum_{j=1}^{p} A_j X_{t-j} + u_{t+h} \tag{12}$$

where z is the shock or instrument for the shock, and X is a vector of control variables. Note that if the shock is identified well, X does not necessarily need to contain all the lagged variables of Y, as they merely act as control variables. Estimation equation 12 has become very widespread in recent years because it is much more flexible than estimating a fully-fledged VAR.

Monetary policy shock Applying the local projection estimation to the familiar three-variable economy leads to the following results. When we use the VAR-based S to identify the shock, the first p (=12) periods closely follow the VAR estimates, and consequently, we have a price puzzle. Thereafter, the output response is considerably larger, and for a brief period even statistically significantly different from the VAR estimates. As is common in local projections, there is considerably more noise in the conditional forecasts.

One weakness of local projections is highlighted when using Gertler and Karadi (2015)'s or Miranda-Agrippino (2016)'s monetary policy shock as z when estimating Equation 12. The fact that the measure in the latter paper is only available from 1990 to 2009 effectively cuts the sample in half (which is not the case in the Proxy-VAR, see above). Additionally, shocks are small, and the local projections have difficulty filtering out the signal from the noise.[10] Inference becomes challenging, in particular in the case of the inflation response. On the other hand, maybe these estimates are "more honest" and the dynamic restrictions of the VAR simply mask how little we know about the effect of monetary policy on inflation, which is a complex animal determined by many current and forward-looking factors. Sometimes, you find smoothed impulse responses where the researchers

Figure 6: Linear local projections: Recursively identified

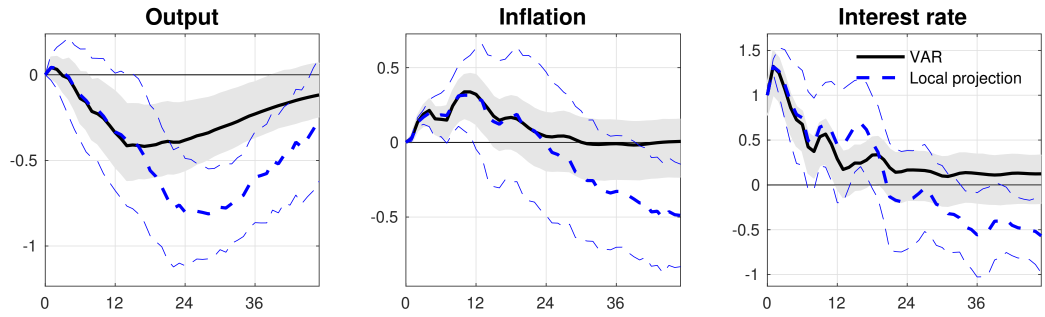

Figure 7: Linear local projections: High-frequency instrument

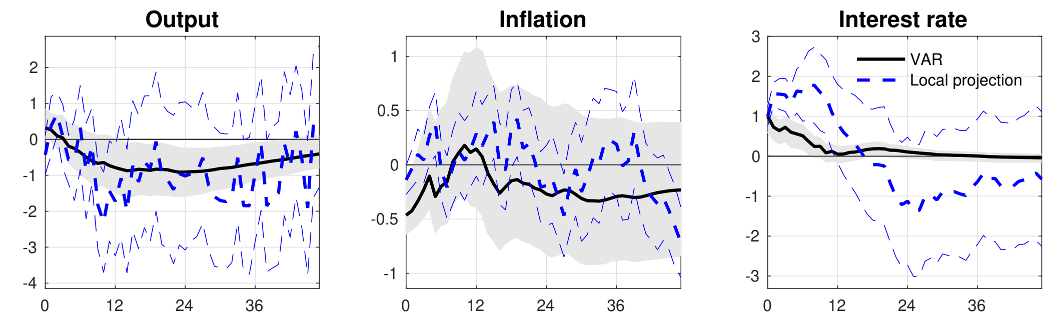

compute moving averages of the mean estimate and upper and lower bounds to make IRFs look smoother. Overall, there are many applications in the literature in which this external identification approach has been used in combination with linear projections, particularly with panel data (because the cross-sectional dimensions increases the econometric power on the one hand and would greatly complicate running a VAR on the other). See e.g. Auerbach and Gorodnichenko (2012b) for an early application, but there are many recent ones that study monetary policy shocks with household (Holm et al., forthcoming) or firm-level (Ottonello and Winberry, 2020) data.

Confidence intervals Since we directly estimate the impulse response in one coefficient, we can use the standard errors on $\hat{\gamma}_h$ to obtain confidence intervals. The only caveat is that the u's will be correlated across h, which is why we need to adjust them for heteroskedasticity and autocorrelation (HAC).

Pros:

- Estimate IRF point-by-point instead of extrapolating period-by-period response implied by coefficients into the too distant future is a more "honest" description of the economy's dynamics
- More robust to misspecification
- Computationally as simple as it gets
- Easy extension to a panel, allowing to exploit cross-section across [pick your unit of observation]

Cons:

- Frequently unstable/"wobbly"
- Only the sample for which all data, including z are available, can be used. If the projection horizon is $h = 48$, we lose 48 additional observations. Inference becomes more challenging.
- Less precisely estimated (for a reason)

6.2 Smooth-transition local projection

Economic theory sometimes implies nonlinearities/asymmetries to be tested/quantified empirically, and sometimes we need empirical evidence to understand the source of state-dependence to refine our theories. In the following nonlinear empirical model, the researcher estimates one of the above LP Equations (11 or 12) but interacts the right-hand side with $(1 - F(s))$, the probability of the economy being in state or regime 1, and once with $F(s)$, the probability of state 2. $F(s)$ is an indicator function further discussed below.

$$Y_{t+h} = (1 - F(s_{t-1}))[\Gamma_{1h} Y_t + \sum_{j=1}^{p} A_{1j} X_{t-j}] + F(s_{t-1})[\Gamma_{2h} Y_t + \sum_{j=1}^{p} A_{2j} X_{t-j}] + u_t \tag{13}$$

Instead of one Γ_h matrix, we now have two, representing the IRFs for each of the regimes. Both Γ_{1h} and Γ_{2h} will have to be identified if no instrument is readily available, just like in the linear case (Section 6.1). To that end, the reader is referred to the concept of conditionally linear Cholesky matrices S_1 and S_2 in Section 6.4.

The state variable F(s) Let us conjecture that an observable (stationary) variable s describes a state of the economy which our impulse response functions $\gamma_{1/2,h}$ depend on. This could be, for example, expansions vs. recessions. In this case, we could use the already familiar output gap series we have used for real output y as the observable since high values indicate booms and low values busts. Call them regimes 1 and 2, respectively. In most applications of Equation 13 we want the interaction to be bounded between 0 and 1, but the output gap is not. In that case, we can apply the logistic function F, which transforms the underlying series into a continuous variable between 0 an 1, giving us a business cycle interpretation: How likely or "how much" is the economy in a recession (state 2) at each point in time?

$$F(s_t) = \frac{e^{-\gamma \hat{s}_t}}{1 + e^{-\gamma \hat{s}_t}}, \quad \gamma > 0, \quad \hat{s}_t = \frac{s_t - \mu}{\sigma_s} \tag{14}$$

We need to make 2 choices in the logistic function: First, the parameter γ rules how "smooth" the transition between states is. Low (absolute) values indicate that extreme regime switches are rare and that the economy frequently hovers around somewhere in-between. The more you raise it, the more sudden we expect changes to be, and the economy is either in a clear boom or bust most of the time. Observe the black and blue lines in Figure 8 illustrating the difference. (When we choose negative values for γ, it implies that low values of s are associated with state 1.[11])

The second one is that before s enters Equation 14, we make sure that it has unit variance. Depending on the look of the underlying series, one should also subtract a constant μ to demean. This constant will affect the way regimes look: Subtracting the mean will let the economy be in either state roughly equal amounts of time. In fact, the value of the percentile chosen is where the value function F will be 0.5. The red line in Figure 8 indicates what happens when we let μ be the 5^{th} percentile of the original series: The transformation only indicates a recession if there is an overwhelming signal for it in the data.

Figure 8: State/regime: Recession probability F(s) with different parameters

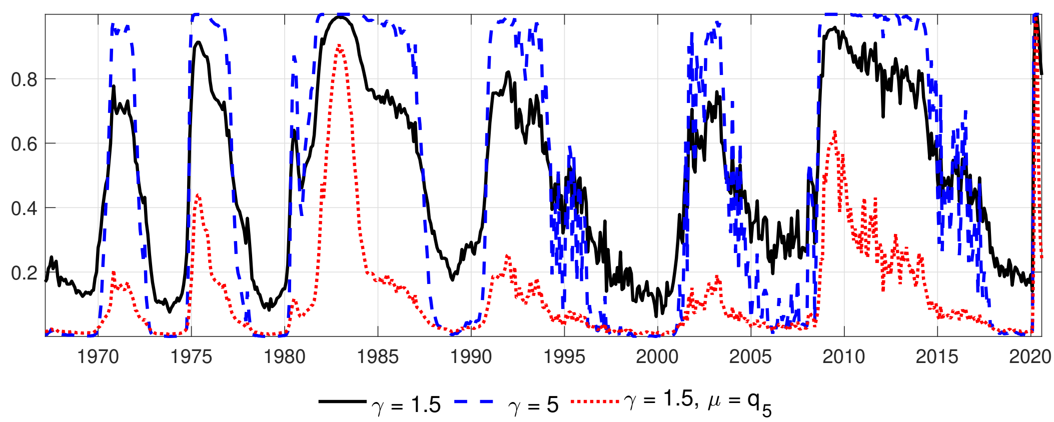

Depending on the (suspected) state-dependence, different functional forms or parameters can be chosen, even more-dimensional ones. The coefficients estimated for each state will show the impulse responses if the economy is fully in that state (while using the information on when we are somewhere in between the extremes). If the state is itself an endogenous variable, which is the case most of the time, the indicator should be lagged by one period in order to ensure we are not assigning the economy to a state where it only is *due to* the shock.

Monetary policy shock Instead of applying the recession/expansion indicator from Figure 8 in this application, let us put that on hold and use another example to show the great flexibility of smooth-transition local projections. Berger et al. (forthcoming) have put forward micro evidence and a theory for the fact that expansionary monetary policy is more powerful if the central bank has been easing monetary conditions for a while (but not too long. The reasons are explained in the very beginning of this handbook), in contrast to when interest rates have been increasing for a while. To test this in a relatively simple exercise, we can define $s_t \equiv i_t - i_{t-12}$ and $F(s_t) = 1$ if $s_t > 0$ and 0 otherwise. In that case $\hat{\Gamma_{1h}}$, or the IRF for regime 1, will give us the economy's reaction for the case where interest rates have fallen in the 12 months prior to the shock. In Figure 9, this is referred to as the easing phase of the interest rate cycle. For $F(s) = 1$, or the economy being in regime 2 where interest rates have been creeping up, the IRFs are depicted in blue.

The main result is that the medium-run effects of monetary policy on output are stronger if the interest rates have already been lowered in the 12 months prior to the shock. This is in line with the hypothesis by Berger et al. (forthcoming). The short-run effects, however, are stronger when the interest rate was on an upward-trajectory to begin with. There is little path-dependence of the inflation effects. Notice also that the behavior of the interest rate itself is quite different: Shocks in the easing phase of a cycle are more persistent, which could in itself be an explanation for the different effects on output.

Figure 9: Nonlinear local projections

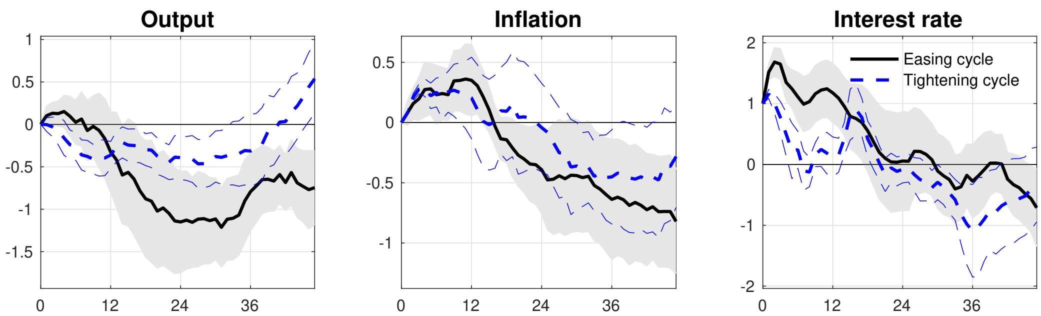

Nonlinear empirical models can address a whole host of questions, interesting for policy or informing more structural models, and they have become extremely popular in recent years. We will look at two more further below, but an overview of a few more papers: Tenreyro and Thwaites (2016) ask if monetary policy is more/less powerful in recessions than expansions. Other examples could be: Are real effects of monetary shocks smaller when inflation is high (because, like in a menu cost model, prices will be more flexible, Ascari and Haber (forthcoming))? Are larger shocks disproportionately effective? Is the transmission channel altered if the zero lower bound is binding (Iwata and Wu, 2006), or if the level of debt (Luigi and Huber, 2018) is high? Are monetary policy effects stronger if announced at different times of the year (because of a seasonality in price stickiness, Olivei and Tenreyro (2007))? Non-linear local projections are a relatively flexible, parsimonious and "easy" (in particular if you have a good instrument) way to study these nonlinearities and have thus become extremely popular in recent years.

Pros:

- All pros for the linear local projection apply.
- Can flexibly accomodate a whole host of nonlinearities, including those difficult to capture in VARs (e.g. size and sign of shock).
- Analytical inference/test of state-dependence is straightforward to implement.

Cons:

- All cons for the linear local projection apply.

6.3 Linear VAR

The linear VAR is already discussed in Section 5. This handbook introduces two nonlinear extensions of this which require *relatively few* (but still quite some) computational adjustments/complications: the smooth-transition VAR (or STVAR, close in spirit to smooth-transition local projections) and the Interacted VAR (IVAR). Both allow for state-dependence, i.e. for the economy (or the vector of endogenous variables Y to respond differently to a shock depending on the values of some other variable. But there are some notable differences in the way those IRFs are constructed and in the way the economy reacts on impact. When we talk about nonlinearities in this seminar, we talk about classes of models (such as the ones below) that are nonlinear in variables but linear in parameters. Such models are more difficult to estimate because of overparameterization: As we will see, the number of parameters increases for the same amount of (typically very limited) data, which can cause problems.

6.4 Smooth-transition VAR

The STVAR is, if you will, a linear combination of the VAR introduced in Section 5.1 and the nonlinearity introduced to local projections in Section 6.2.

$$Y_t = (1 - F(s_{t-1}))[\sum_{j=1}^{p} A_{1j} Y_{t-j}\,] + F(s_{t-1})[\sum_{j=1}^{p} A_{2j} Y_{t-j}\,] + u_t \tag{15}$$

$$E(u_t) = 0, \qquad E(u_t u_t') = \Omega_t$$

$$\Omega_t = (1 - F(s_{t-1}))\Omega_1 + F(s_{t-1})\Omega_2 \tag{16}$$

$$F(s_t) = \frac{e^{-\gamma \hat{s}_t}}{1 + e^{-\gamma \hat{s}_t}}, \quad \gamma > 0, \quad \hat{s}_t = \frac{s_t - \mu}{\sigma_s}$$

$F(s)$ can be interpreted as the probability of the underlying regime 2. The user is referred to a discussion of the logistic transformation in Section 6.2. Equations 15 and 16 show that the nonlinear model is set up to be nothing but a weighted sum of two linear models: a model for regime 1 with the estimated coefficients for the lagged variables, A_1, as well as the covariances of residuals in Ω_1, and likewise for regime 2.

If we have estimated matrices of parameters A_1 and A_2, we can recover the implied impulse responses for each regime exactly like in the linear model. In computing these conditionally linear impulse response functions, we implicitly assume that the economy stays in the same regime as it was at the time of the shock. That means that IRFs can look explosive, i.e., do not revert back to zero/trend. That is not problematic because in each horizon h, there is a possibility of switching to the other regime. Another consequence is that the size and sign of the shock do not matter (because we identify them separately on Ω_1 and Ω_2, respectively), so there is no nonlinearity in that sense.

Computational issues There are two main computational hurdles to take, relative to the linear VAR. First, we cannot estimate the above equations by OLS and we therefore need to rely on numerical methods instead. Similar to how the maximum likelihood is explained in Lütkepohl and Netšunajev (2014), one can stack the relevant parameters to be estimated into one vector $\theta = \{A_1, A_2, \Omega_1, \Omega_2\}$ and choose them such that they maximize the log-likelihood function:

$$\log L(\theta) = \text{constant} - (t/2)\log|\Omega_t| - (1/2)\sum_{1}^{t} u_t' \Omega_t^{-1} u_t$$

However, the second issue still applies: For the bootstrapping loop, we typically draw from the errors u and simulate the variables in Y given A. In this case, however, Y also depends on $F(s)$, which is a non-trivial function of a variable related but not necessarily included in the vector Y. Therefore, most applications rely on Bayesian estimation techniques as available in the online appendix to Auerbach and Gorodnichenko (2012a) and

Auerbach and Gorodnichenko (2013).[12] Markov chain Monte Carlo (MCMC) techniques have the advantage that they have built-in parameter uncertainty, which allows for constructing errors bands without relying on bootstrapping.

Monetary policy shock This paragraph illustrates the (dis-)advantages of the STVAR with a simple application to a monetary policy shock. Let us assess if the effects of monetary policy are affected by the credit cycle. Credit, and in particular mortgages, play an important role in the monetary transmission mechanism, because changes in the interest rate directly affect the cost of credit and therefore, indirectly, housing (Iacoviello, 2005). The regime variable s is defined as the y/y growth rate of the ratio of nominal mortgages to households and nominal GDP. The series on mortgages is a quarterly one, so the entire model is estimated at the quarterly frequency. The details and how we transform the series in the indicator $F(s)$ is discussed in the video tutorial (see Section 7). With respect to the choice of s, it is usually required that this series is stationary - one could not use the mortgage/GDP ratio as a level, which has a clear upward trend. At the same time, it should not be too unstable, i.e. switch between regimes from period to period. One can try to reduce the number of parameters by not including non-essential series and/or reducing the number of lags. Figure 10 shows the impulse responses for output and inflation to a 1% increase of the monetary policy rate when a.) mortgages have grown faster than GDP during the four quarters preceding the shock and b.) credit growth has been slower than output growth (and – crucial for interpretation – the economy stays in the "low credit growth" regime).

Figure 10: Smooth-transition VAR

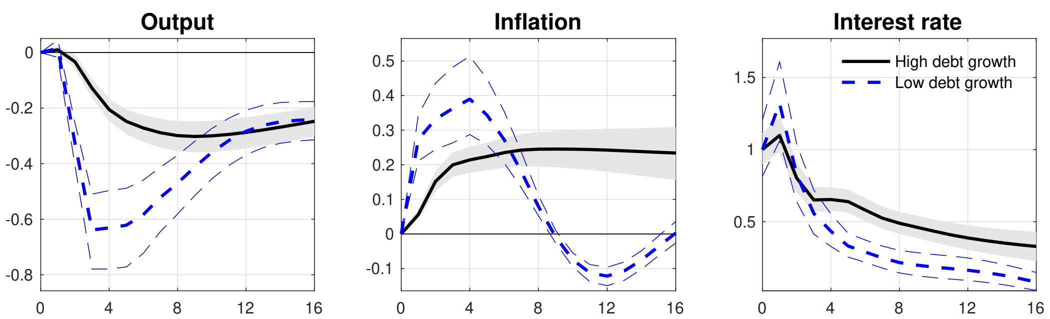

According to these results, the effects on output are more contractionary if mortgages have been growing slowly. As discussed above, the fact that some IRFs are not converging back to zero is not a problem, as in reality the economy will switch to the other regime (exogenously or endogenously), so we will never observe these explosive movements in the data. How could we give these results an economic interpretation? Ex ante, one might have expected the opposite. In times where the economy is highly leveraged – many households have large mortgages – an increase in the interest rate affects more people and creates larger demand effects. However, the results here indicate the opposite. A potential explanation is that most households have long-term fixed-rate mortgages. If many households have obtained a new mortgage in the past, it means they have "locked in" the level of the interest rate on their mortgage. Therefore, monetary policy affects housing and borrowing choices at the margin of "new" borrowers: If credit growth has been low, there are many households without a mortgage or with the potential to increase their mortgage (for a renovation or consumption). This decision is very responsive to monetary policy, leading to these large demand effects. If everyone's houses are highly levered already (with a fixed interest rate), increasing the interest rate on *new* mortgages does not have as large effects.

Three general remarks of caution: First, the above explanation is *one* hypothesis. Robustness tests and the inclusion of other variables (in particular mortgages, ideally new and existing) would have to confirm the story. Second, the underlying nonlinearity is *one* dimension of many that are possibly correlated. If the low debt growth regime coincides with recessions, we might pick up an effects that comes from the fact that we have a recession, rather than low credit growth per se. To make a definitive statement, we would have to address this one way or another. Third, the error bands in Figure 10 are implausibly close. This is a (bad) feature, not a bug, of the STVAR whenever we identify the shock on a variable that is order "late" in the system. The toolbox of codes includes the replication of Auerbach and Gorodnichenko (2012a), where the shocked variable is ordered first, which is when the STVAR really comes to fruition. Alternatively, the codes are organized in a way that splits the estimation from the identification part, so one could potentially apply the IV way of identifying the VAR (Section 5.3) to the reduced-form residuals for each posterior draw, which would likely make the confidence bands look much better because there are no zero restrictions.

Pros:

- Observed values are a linear combination of two data-generating processes $\rightarrow$ elegant and serious treatment of state-dependence
- IRF is state-dependent even on impact (because of two estimated Ω's)

Cons:

- Observed values are a *linear* combination of *two* data-generating processes $\rightarrow$ overly restrictive and difficult to estimate on realistic samples
- Number of parameters grows exponentially
- Computationally demanding

6.5 Interacted VAR

Finally, we are going to look at the effects of monetary policy in expansions vs. recessions. Tenreyro and Thwaites (2016) found that the real effects a central bank shock are larger in expansions. This would be bad news: Monetary policy seems to be least effective when we most need it. In the application here, the underlying state is the output gap itself (which part of the endogenous variables). As with many "states", the output gap contains information that is far beyond the binary extremes (boom vs. recession), and the IVAR can make use of this.

The Interacted VAR is estimated as a linear model with an interaction term, which makes it "easier" econometrically and it requires fewer parameters to be estimated. We regress the values of the variables in Y not only on their lags but also on their interaction. In the simplest version discussed here, we choose two variables y_1 and y_2, both of which are themselves elements of Y (e.g. the output gap, or again the interest rate).

$$Y_t = \sum_{j=1}^{p} A_j Y_{t-j} + \sum_{j=1}^{p} B_j [\, y_{1,t-j} \times y_{2,t-j}\,] + u_t \qquad E(u_t) = 0,\, E(u_t u_t') = \Omega \tag{17}$$

Why is this a nonlinear model? If a shock realizes in period t and has an immediate effect on y_{1t}, then $y_{1,t+1}$ will also depend on the value of y_{2t}, which is itself affected by the shock. The interaction term should therefore include both the variable we believe to be shocked (e.g. the interest rate) and the variable we believe to describe the "state" (e.g. the output gap). The advantage of this is that we can use the familiar OLS function to estimate this model like a linear VAR with exogenous variables, which simplifies matters a lot.

However, tractability and parsimony come at a cost: Because there is only one estimated Ω, one restricts impulse responses to be the same across states in the period of the shock (regardless of the ordering of variables). In effect, we set up a nonlinear model saying that there is no nonlinearity on impact, which is the biggest conceptual flaw of the IVAR. However, in the case where the shocked variable is ordered last, this is no issue, because all variables are restricted to have no movement on impact anyway if we do a Cholesky decomposition. The second drawback is that both y_1 and y_2 have to be part of the endogenous variables, which can (but does not have to) be restrictive.[13]

Impulse responses function In this model, we cannot compute impulse responses the traditional way, i.e., we have no closed-form solution. Instead, we need to rely on a numerical method. The idea of generalized impulse response functions or GIRFs (Koop et al., 1996) is the following:

1. We take a particular initial condition, which is a single row of actual observations of all variables and the relevant lags $\omega_{t-1} = \{Y_{t-1}, ..., Y_{t-p}\}$. We call this a "history" because it describes one particular historical realization of the data (and its interactions).
2. We simulate a theoretical path all the variables would take in the subsequent h periods absent of any shocks, using the estimated matrices A and B.

3. We simulate the same path with the same coefficients but shock the system with ε at period t. ε_t is drawn from the (orthogonalized) residuals we estimated. The impulse response function is then the difference between the two paths:

$$GIRF_{Yt}(h, z_t, \omega_{t-1}) = E[\,Y_{t+h}\,|\,\varepsilon_t, \omega_{t-1}\,] - E[\,Y_{t+h}\,|\,\omega_{t-1}\,]$$

4. Repeat 2 and 3 multiple (e.g. 500) times, drawing new shocks and in the end average impulse responses across draws. (Sometimes, the outlier-robust median yields better results.)
5. We repeat 2 to 4 for all histories, i.e. each month or quarter of the actual data and compute the average of all histories (in the linear case) or of certain selected histories belonging to what we believe is a particular regime (in the nonlinear case).

GIRFs are in principle applicable to all linear and nonlinear VARs, including the STVAR (see e.g. Caggiano et al. (2015)). We can also bootstrap to obtain standard errors just like in the earlier models. GIRFs can be computationally demanding as the processing time increases linearly in the amount of initial conditions/time periods in the data, but we can do it all with frequentist methods. (Parallelization of this process is possible.) Another drawback of GIRFs is that there is nothing that prevents either or both of the simulated paths to be explosive, and we need to disregard those draws, increasing the amount of necessary iterations. Extra caution should be applied to models with a high share of explosive paths. This is particularly problematic if the nonlinearity is somewhat of a corner case – for example, if you are trying to distinguish between histories at the zero lower bound vs. higher interest rates. The following application, however, is not such a case.

Monetary policy shock In order to assess the effects of monetary policy in booms and busts, we make a number of changes. First, we again use quarterly data because it is easier to achieve stable dynamic coefficients, and the GIRFs are less forgiving in that respect. The interaction variables y_1 and y_2 are the output gap and the interest rate. Instead of the usual Federal funds rate, we take the so-called shadow-rate by Wu and Xia (2016). This is an imputed interest rate that incorporates the fact that during the recovery from the Great Recession, unconventional monetary policies made the actual policy stance lower (indeed negative) than the observed Fed funds rate, which was bound by the nominal value of zero.[14] Finally, we need to group the histories. To do so, we take all values of y_1 and define as recessions those quarters (histories) where its value was below the 12th percentile, which is about the fraction of quarters the U.S. economy spends in recessions according to the NBER. The respective impulse response functions are depicted in Figure 11.

Figure 11: IVAR

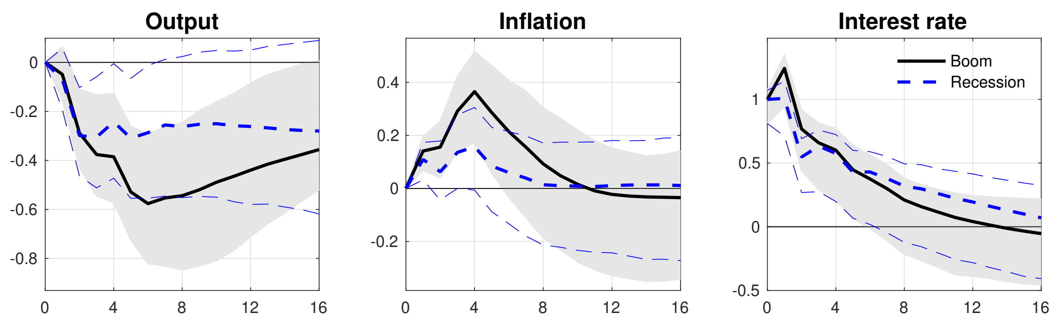

First, observe that the impulse responses have a similar shape and magnitude than with data at the monthly frequency: A 1% interest rate hike lets output contract by about 0.5% and we have quite a price puzzle. Comparing the dashed blue and black solid line, we see that the output contraction is only about half as strong for a monetary policy shock happening in a recession, and is not statistically significant after 2 years (8 quarters). At the same time, however, it is questionable to conclude that the output response is statistically significantly different in booms vs. recessions.

A second notable feature is that contrary to the STVAR, impulse responses tend to mean-revert. When computing GIRFS, we only condition the economy to be in the respective regime on impact. Afterwards, the economy will "switch" back and forth with the usual dynamics (and also affected by the shock). Since the economy is in expectation a stable system, impulse responses mean-revert. This is not necessarily the case in the STVAR, where the IRFs stipulate what would happen if the economy stayed in the respective regime the entire time period – a behavior we will never observe because the economy does move to different regimes. Another typical finding with GIRFs is that the mean estimates are not necessarily centered within the 90% bootstrapped confidence bands in the IVAR. If only a handful of histories/initial conditions or draws of the residuals drive the shape of the IRF, then those carry a lot of weight when averaging in steps 4 and 5 of the GIRF simulation. The IVAR is less frequently used in the literature than the STVAR (which has been hyped for a number of years), but can be relatively flexibly applied to many nonlinear contexts such as uncertainty shocks at the zero lower bound (Caggiano et al., 2017).

Pros:

- Parsimonious way to implement nonlinearities, e.g. only requires the estimation of one variance-covariance matrix and can thus be estimated with OLS.
- The model does not impose nonlinearity, i.e., if the data suggests that there is none, the estimated parameters in B will tell us. Additionally, we acknowledge that the nonlinearity is not inherently binary such as in the STVAR.
- GIRFs can be different for shocks of different times ("state"), but also sign and size.

Cons:

- Because we only estimate one Ω, we essentially restrict the shock response on impact to be the same regardless of the state. Thus, IRFs can only show state-dependence in the medium-run.
- Computing GIRFs can, for some histories, lead to explosive paths.
- Computationally more intense as we have to draw and simulate many times.
- The "state" variable y_2 has to be part of the vector of endogenous variables Y, and it has to be stationary.

6.6 Estimated DSGE model

Macroeconomists frequently set up models that represent structural relationships we believe to be true in the economy. We can bring these structural relationships to the data and estimate their parameters given the observable comovement of the variables in the model. As a by-product, we are identifying the structural shocks and can simulate the economy's response to them.

Structure Take the following example of the Euler equation: In many macroeconomic models, forward-looking households decide how much to consume and work today, subject to a budget constraint. They maximize

$$E_0 \sum_{t=0}^{\infty} \beta^t \left(\frac{C_t^{1-\sigma}}{1-\sigma} - \frac{N_t^{1+\psi}}{1+\psi} \right) \qquad s.t. \quad P_t C_t + B_t \leq (1+i_t) B_{t-1} + W_t N_t$$

In words, this translates to: Maximize the expected discounted utility streams from consumption, minus the disutility from having to work. At the same time, the nominal flow budget constraint must always hold. The outflow of money (for nominal consumption PC or for new investments in bonds B) cannot exceed the inflow of money (from previous investments B_{t-1} or from working WN). If we set up the Lagrangian and derive the first-order conditions with respect to consumption C and investments B, we can go on to derive the Euler equation 18.

$$\frac{\partial \mathcal{L}}{\partial C_t} = C_t^{-\sigma} + \lambda_t P_t = 0 \Rightarrow \lambda_t = -\frac{C_t^{-\sigma}}{P_t}$$
$$\frac{\partial \mathcal{L}}{\partial B_t} = \lambda_t - \beta E_t\{\lambda_{t+1}\}(1+i_t) = 0$$
$$\Rightarrow \frac{C_t^{-\sigma}}{P_t} = \beta E_t\Big\{\frac{C_{t+1}^{-\sigma}}{P_{t+1}}\Big\}(1+i_t)$$

$$C_t^{-\sigma} = E_t\{C_{t+1}^{-\sigma}\}\frac{P_t}{E_t\{P_{t+1}\}}\beta(1+i_t)$$
$$C_t = E_t\{C_{t+1}\}\left(\frac{P_t}{E_t\{P_{t+1}\}}\beta(1+i_t)\right)^{-\frac{1}{\sigma}} \tag{18}$$

In words, this equation says the following: All else equal, if we expect prices to increase tomorrow ($E_t\{P_{t+1}\} \uparrow$, we will consume more today. (The $-1/\sigma$ in the exponent flips numerator and denominator, so higher price expectations mean the right-hand side as a whole increases.) If $\beta \uparrow$, it means we become more patient, the RHS decreases, and our current consumption decreases because we postpone it to the future. If the interest rate goes up, we decide to forego some of the consumption today to make use of higher returns on savings and thus C_t will also be low.

This will give us one dynamic relationship between (expected) inflation, interest rates and real consumption. Imposing this relationship on the data will not only potentially allow us to estimate the parameters σ and β, but also consider any deviations from these equations as structural shocks, in this case here, a demand shock, i.e., a change in consumption that is neither driven by inflation expectations, or interest rates.

The combination of possible equations come in many shapes and forms. The one that is considered here is the simplest version of a New Keynesian model meaningful for the analysis of monetary policy. It is only briefly summarized here, but the exact mechanics and derivations can be found in chapter 3 of the Galí (2015) textbook.

$$y_t = E_t\{y_{t+1}\} - \tfrac{1}{\sigma}(i_t - E_t\{\pi_{t+1}\}) \tag{19}$$
$$\pi_t = \beta E_t\{\pi_{t+1}\} + \kappa\, y_t \tag{20}$$
$$i_t = \phi_y y_t + \phi_\pi \pi_t + z_t^i \tag{21}$$

where $\kappa \equiv \frac{(1-\omega)(1-\beta\omega)}{\omega}(\sigma + \frac{(\psi+\alpha)}{1-\alpha})(\frac{1-\alpha}{1-\alpha+\alpha\epsilon})$. The first one is a version of the Euler equation 18, however, with a certain amount of changes. First, in this simplest closed economy we assume that consumption is the only form of demand (there is no government, and no capital to invest in); therefore, consumption is also GDP. Second, we have log-linearized[15] the equation around the steady state.[16] Equation 19 is referred to as the dynamic IS equation. Equation 20 – called the New-Keynesian Phillips Curve – can be derived from first principles and says the following: Inflation is a weighted product (or in logs: sum) of expected inflation and of the output variable. The second component with κ is interesting: It gives us the slope of the Phillips curve which is most prominently (negatively) dependent on ω, which is the parameter ruling how sticky the price-setting mechanism of firms is. (It also depends on the elasticity of demand ε and the output elasticity of capital α.) If prices are flexible, then an overheating economy ($y \uparrow$) quickly leads to higher inflation. If prices are sticky, then the economy needs to move a lot in order to have an effect on the aggregate price level, and the NKPC is relatively flat.[17] To close the model, we pick a third equation that is relatively ad hoc. It simply says that the central bank will steer the interest rate such that it systematically responds to both output and inflation, but it can also deviate from this so-called Taylor rule by the variable z^i.

VAR representation Let us show that the 3-equation model actually has a representation that looks something like a VAR, i.e., where all three variables are expressed in terms of each other's lags. (This is not a solution method, simply an illustration). Rewrite equations 19 to 21 in state-space form:

$$\begin{pmatrix} -1 & -\frac{1}{\sigma} & 0 \\ 0 & -\beta & 0 \\ \phi_y & \phi_\pi & -1 \end{pmatrix} \times \begin{pmatrix} E_t\{y_{t+1}\} \\ E_t\{\pi_{t+1}\} \\ E_t\{i_{t+1}\} \end{pmatrix} = \begin{pmatrix} -1 & 0 & -\frac{1}{\sigma} \\ \kappa & -1 & 0 \\ 0 & 0 & 0 \end{pmatrix} \times \begin{pmatrix} y_t \\ \pi_t \\ i_t \end{pmatrix} + \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \times \begin{pmatrix} 0 \\ 0 \\ z^i_{t+1} \end{pmatrix}$$

$$G \times Y_{t+1} = H \times Y_t + I \times \varepsilon_{t+1}$$

$$Y_t = AY_{t-1} + S\varepsilon_t$$

This is our very first Equation 1. The last step, however, is only possible if G is invertible, with $A \equiv G^{-1}H$ and $S \equiv G^{-1}$. If we can re-write the 3-equation New Keynesian Model as a VAR, that would mean the parameters we estimate are informative of the structural parameters of the model.

Shocks It should be clear at this stage why this system of equations is called a *Dynamic* Stochastic *General Equilibrium* model. But we are missing the *Stochastic* part. Of course, the three equations alone will never be able to perfectly fit the data. DSGE modellers typically impose further structure onto these residuals, which is why we refer to them as shock processes z. Let each of the three residuals follow an AR(1) process like

$$z^y_t = \rho_y z^y_{t-1} + \varepsilon^y_t, \qquad \varepsilon^y_t \sim \mathcal{N}(0, \sigma^2_{\varepsilon^y}) \tag{22}$$

$$z^\pi_t = \rho_\pi z^\pi_{t-1} + \varepsilon^\pi_t, \qquad \varepsilon^\pi_t \sim \mathcal{N}(0, \sigma^2_{\varepsilon^\pi}) \tag{23}$$

$$z^i_t = \rho_i z^i_{t-1} + \varepsilon^i_t, \qquad \varepsilon^i_t \sim \mathcal{N}(0, \sigma^2_{\varepsilon^i}) \tag{24}$$

We can call ε^y_t a normally distributed demand shock. Both the persistence of the shock process ρ_y and the variance of the innovation $\sigma^2_{\varepsilon^y}$ are parameters of the model that will be estimated. We further refer to ε^π_t as a cost-push shock and ε^i_t as the familiar monetary policy shock. According to equation 19, a surprise increase in of ε^i_t will introduce an immediate decrease in GDP which will also transmit to a lower level of inflation according to equation 20. (In fact, it is the other way around: Since prices are sticky, they do not fully adjust to the new nominal interest rate, and thus the *real* interest rate also moves, leading to a shift in the dynamic allocation of consumption by households). A notable shock is absent from the analysis at this point, namely the technology shock.[18]

Estimation Estimating even simple DSGE models is, as you probably can assess at this point, a major endeavor. Luckily, there is a publicly available toolbox of Matlab codes called *Dynare* that gives us the chance for a first shot at an estimated DSGE model. Dynare allows both to simulate DSGE models (to calculate impulse response functions given the parameters the model is calibrated with), or to estimate them. In the latter case, the algorithm needs to know the structural equations (as derived above in deviation from steady state), the data series, and *prior guesses* of the parameters.

Regarding the data series, notice that we need to define (at least) as many data series as there are structural shocks in the model. Although this constraint is normally non-binding, it can be with larger models where not all variables have a directly observable counterpart. In our case, it is relatively straightforward to use the output gap, inflation and interest rate series we have already used prior.

Dynare uses Bayesian estimation techniques, which per se go beyond the scope of this course. Apart from the data and the model itself, the user needs to specify prior beliefs about the parameters she wants to estimate. By parameters, we mean both the structural parameters such as β or the policy reaction function ϕ_y and ϕ_π, but also for the shock processes themselves, including the variances of the structural shocks. The full vector of parameters of our very small model would be $\theta = \{\sigma, \beta, \psi, \omega, \phi_y, \phi_\pi, \rho_y, \rho_\pi, \rho_i, \sigma_{\varepsilon^y}, \sigma_{\varepsilon^\pi}, \sigma_{\varepsilon^i}\}$. Some of them, like the discount factor, we can calibrate with relative certainty. Others are less certain, but we can tell the optimizers how "tight" we want our priors to be. The model will try and choose the parameter vector such that they maximize the likelihood of generating the data series we submit to the program.

These priors can have one of the five following shapes:

Table 2: Prior shapes, to be specified in Dynare

Name:	Density moments: Distribution(moments)	Range
normal_pdf	$N(\mu, \sigma)$	$\mathbb{R}$
gamma_pdf	$G_2(\mu, \sigma, p_3)$	$[p_3, +\infty)$
beta_pdf	$B(\mu, \sigma, p_3, p_4)$	$[p_3, p_4]$
inv_gamma_pdf	$IG_1(\mu, \sigma)$	$\mathbb{R}^+$
uniform_pdf	$U(p_3, p_4)$	$[p_3, p_4]$

There are a maximum of four prior moments we can specify: the mean μ, the standard deviation σ, and the third and fourth moments p_3 (by default 0) and p_4 (by default 1). If we knew that a parameter lies between 0 and 1 (like the probability of microeconomic price flexibility) we could submit "omega, uniform_pdf, , , 0, 1;" or even better "omega, beta_pdf, 0.7, 0.05". If we know that a parameter has to be positive, the convention is to use the inverted Gamma distribution with a small mean and large variance. The command for the prior of the monetary policy shock size σ_ε^i could for example be "epsi, inv_gamma_pdf, 0.1, 2".

Monetary policy shock The online toolbox includes Matlab, data and .mod (Dynare) files in order to estimate a 3-equation New Keynesian model. Once we have the estimated parameters, we also know the implied structural shocks at each point in time. The stochastic simulator then takes the estimated parameters (posterior modes) and simulates data in order to compute the impulse response functions. In the case of our very imperfect model, they look like the black solid line depicted in Figure 12.

Figure 12: New Keynesian DSGE models

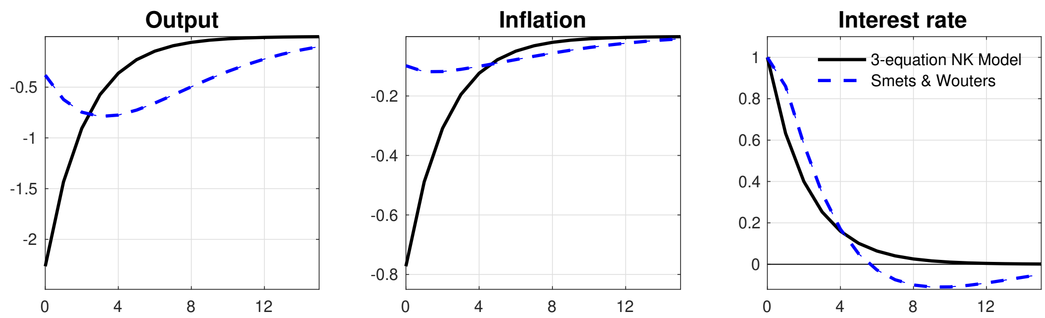

The model finds that a 1% surprise increase leads to an immediate contraction of the output gap by 2 percentage points, which is considerably larger than in any of the empirical exercises – Remember that when using Cholesky, we actually *impose* a zero response on impact. At the same time, the effect dies out within a matter of months. Almost all our statistical approaches found that the peak impact on economic output is reached after a year or two. Monetary policy operates with a lag. The 3-equation New Keynesian Model is clearly at odds with the hump-shaped pattern suggested by the data.

Pros:

- Makes use of some of the fundamental relationships of economic variables we "know" about – More (micro-founded) Economics, less Statistics.
- Allow for better explanations concerning the sources of economic fluctuations compared to VARs; address the non-fundamentalness problem VARs are criticized for

Cons:

- Structural equations and rational expectations hide the (huge) degree of uncertainty of even the most basic economic theory.
- Weak forecasting performance
- Technical demands

6.7 Extensions

To get more realistic impulse responses, we need a richer model. The online toolbox also contains replication codes for the Smets and Wouters (2007) model, a 14-equation DSGE model with 7 shocks like it is used in many central banks around the world. Relative to the 3-equation model introduced above, it contains additional blocks for investment and capital utilization, a risk premium, a monopolistically competitive labor market (i.e. a wage markup), and a government with a budget constraint. These additional frictions are crucial in order to get more realistic IRFs, as we find them in the data. The generated impulse responses for the three variables of interest are represented by the dashed blue lines in Figure 12.

DSGE models can incorporate nonlinear dynamics, too, even though it comes with additional computational challenges that are probably most likely too demanding for a seminar paper. Nevertheless, it might become important for future work. An interesting and highly relevant question is how the zero lower bound on nominal interest rates changes the dynamics of the macroeconomy: Because rational agents would not hold money if the interest rate was negative, there is a lower bound at zero (in principle). The model description is relatively simple, as we can simply require i to be non-negative, i.e. add an inequality constraint to the system of equations 19-21 that says $i_t \geq 0$. This constraint does not always hold and is thus called an occasionally binding constraint. Guerrieri and Iacoviello (2015) have developed a Dynare-compatible set of codes that solves (but does not estimate) models with OBCs.

It's less interesting to study the effect of a monetary policy shock at the zero lower bound on output in this case, because an increase in i will have the same effect on the other equations, and a decrease will not be possible. However, it makes the effects of the other shocks state-dependent: If there is a negative shock to demand or a positive shock to supply, both of which reduce inflation, the central bank's Taylor rule would normally suggest to reduce interest rates. If those are constrained, however, monetary policy is "too tight" and that will lower output.

Because the zero lower bound led to major central banks experimenting with "unconventional" monetary policies such as directly purchasing financial assets to provide markets with liquidity, a short discussion of the effects of these quantitative easing (QE) policies is in order. A recent paper by Sims et al. (forthcoming) develops a 4-equation New Keynesian model that incorporates processes for bond holdings by the private sector and the central bank, respectively. Figure 13 shows the impulse responses of a conventional and an unconventional monetary policy shock in their model. All parameters are calibrated, but the size of the Federal Reserve balance sheet in relation to nominal GDP is used to estimate the size of the shocks.

Figure 13: New Keynesian DSGE model with quantitative easing

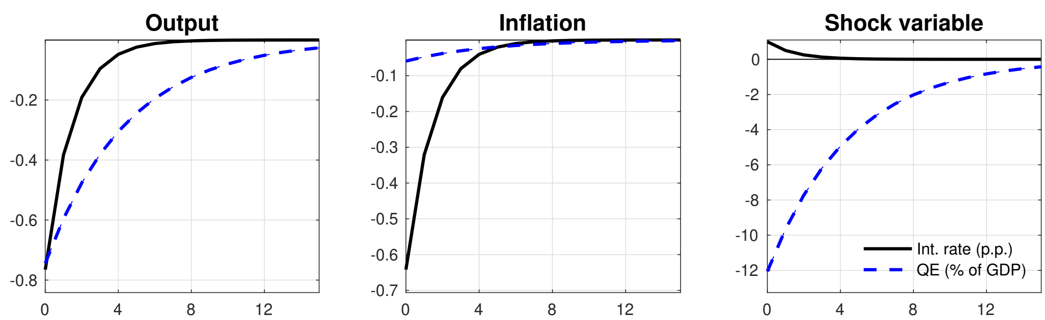

Taking the model wholly seriously, both the increase of the nominal interest rate (through the Euler equation) and the reduction of central bank reserves to withdraw liquidity from the economy (negative QE) decrease GDP, even though the latter needs to be substantial to have a quantitatively meaningful effect. The starkest difference between the two is that relative to the size of the output effect, the QE shock is less deflationary than the conventional monetary policy tightening. Through the lens of the Sims et al. (forthcoming) model, QE shocks bear a cost-push element rationalized as follows: Favorable credit conditions reallocate resources from savers (households) to borrowers (entrepreneurs who do not supply labor). The former increase their labor supply (negative wealth effect), exerting downward pressure on wages and thus marginal cost for the firms and finally inflation. The relative strength of this effect depends on the model parameters, but for the calibration chosen in Figure 13, the inflation response is minuscule even for large QE shocks. Their online appendix also solves the 4-equation model under the occasionally binding constraint of a zero lower bound.

7 Toolbox and tutorials

The "macrometrics" page on my website contains code to the models introduced in this handbook as well as video tutorials walking through some of the particularities in detail. All the approaches to shock identification discussed in Section 5 are present with an application to the monetary policy shock obtained from monthly US data. The code thus also produces the figures depicted throughout this handbook. It also provides a primer on how organize your macroeconomic time series and codes. Some prior knowledge of matrix multiplication and notation in Matlab and/or R is required.

The full list of tutorials is:

1. Organizing data
2. Linear local projections
3. Linear VAR: Cholesky identification
4. Linear VAR: Sign restrictions
5. Linear VAR: High-frequency identification
6. Smooth-transition VAR
7. Interacted VAR

Accessing the tutorials requires a direct link to the online videos. If you do not participate in the seminar and would still like to obtain these links, please get in touch. Reports of mistakes or missing clarifications in the code are greatly appreciated.

8 Other useful tools

This Section presents a handful of useful tools that can be applied to most of the above models and that one might encounter when reviewing the literature. The list is subject to extensions for future iterations of this handbook, and there are no video tutorials available for these extensions as of now.

8.1 Scenario analysis

For policy makers, an intriguing feature of empirical models is that they allow answers to "What if...?" questions. What if the Fed had decided on a different monetary policy stance at a certain meeting? What would have been the expected path of one variable if another variable had behaved differently? Two such scenarios are visualized and discussed below. Both are imperfect and use relatively simple, not state-of-the-art simulations of a series of hypothetical shocks that the user feeds into the model.[19]

Counterfactual example I: A welcome present for President Trump In the first example, we simulate the path of a one-time monetary policy shock. Donald Trump was inaugurated as President in January of 2017. The economy had recovered from the Great Recession, output was close to potential and inflatiion close to target. Therefore, the FOMC, which is the committee making monetary policy decisions, and its Chairwomen Janet Yellen had been slowly raising rates for a year. Assume that the FOMC would like to make incoming Presidents a gift by a one-off expansionary monetary policy shock that lower the interest rate back to zero – Short-sighted politicians tend to like expansionary monetary policy.

Figure 14: Counterfactual example I (expansionary shock in 2017)

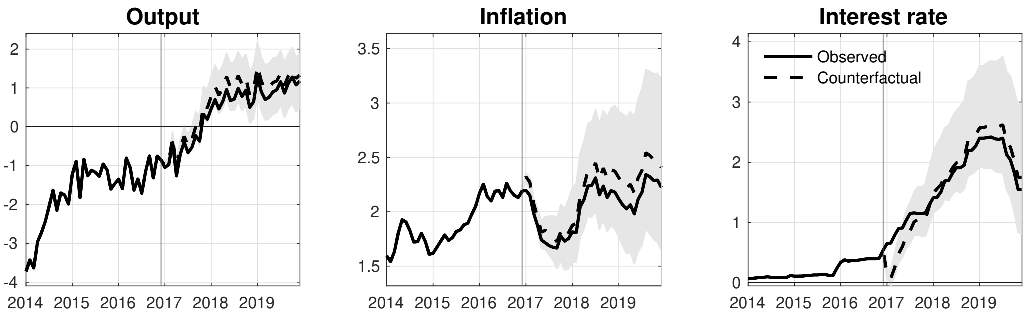

How would output and inflation have evolved in this scenario? According to Figure 14, not very differently. The plots take into account the shocks that hit the economy before and after January 2017, but additionally shock the interest rate with -0.5 in that month. The reason why the simulated paths for output and inflation (dashed lines) closely align with the actual data are two-fold. First, it is established that while "random" monetary policy shocks do have significant effects on output, they actually explain a relatively small part of its fluctuation (8% in Ramey (2016)). The second is that due to the coefficients of A we use to simulate the paths, the post-shock values of output and inflation require the central bank to increase rates again afterwards, and to do so quicker then they did in reality. Therefore, the actual stimulus is only temporary.

Counterfactual example II: The cost of the zero lower bound This is different in the second example. The baseline model is estimated with the Federal funds rate as the interest rate, which in the U.S. never reached negative territory due to concerns of financial stability with negative rates. Instead, the Fed kept rates at or very close to zero between 2009 and 2015. Figure 15 shows the prediction of the VAR under a scenario where the central bank lowered the interest rate to -2% and gradually raised it subsequently. In every period, the sign and size of the monetary policy shock is chosen in order to reach exactly these target values.

Figure 15: Counterfactual example II (absence of zero lower bound)

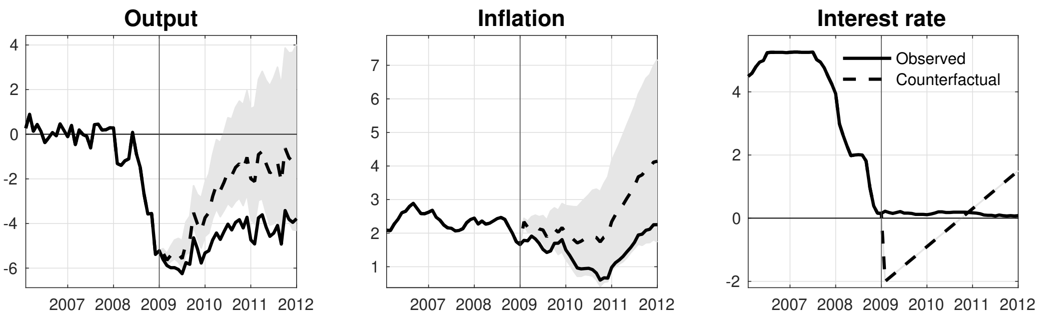

As we see in the left-hand panel of Figure 15, output would have recovered much faster. In fact, the output gap would have been essentially closed in 2011 or 2012, instead of 2015. This is despite a more hawkish stance from 2011 and forward, which would presumably have given the Fed more room to lower rates when confronted with a new crisis – such as Covid-19. We can conclude that not being able or willing to move rates into negative territory early on contributed to a sluggish recovery.

The exercise also illustrates caveats of these types of counterfactuals: First, the confidence bands (66%) are wide because we hit the economy with a series of shocks, the effect of each of which comes with its own uncertainty. Second, our VAR cannot account for (forward-looking) expectations which are often a function of the very policy change we are trying to assess.[20] Third is that in reality, the Fed *did* do something when it hit the zero lower bound in 2009: unconventional policies such as forward guidance (i.e. trying to change expectations) and quantitative easing. Because this is not accounted for by the three variables shown, the model and its shocks are likely badly estimated in this period.

8.2 Misspecification tests

When we estimate Equation 1, we don't know the "true model" of the economy. The reduced-form residuals (u) should be normally distributed and serially uncorrelated and if they are not, it can be an indication that our model is mis-specified. The extra toolbox contains some functions that visualize the u's by means of histograms and QQ plots, select the best number of lags according to the most common information criteria and test for heteroskedasticity. A more comprehensive overview is provided by Kilian and Lütkepohl (2017).

The reader should also keep in mind that these are not black and white issues: a model is not either well or mis-specified. In fact, since VARs are merely a representation of the data-generating process, they are certain to be at least marginally misspecified. Also note that, although these tests are often given a lot of emphasis in courses and textbooks, they are rarely seen in published empirical papers. Misspecification is often inevitable when working with time series models and econometricians frequently ignore it. Remember that the focus of this seminar is on the identification and nonlinearities of macroeconomic shocks and that's what you should primarily spend your time, as well as keystrokes, on.

8.3 How to estimate macroeconomic shocks after March 2020

The Covid-19 pandemic and the resulting economic crisis led to unprecedented variation in several macroeconomic variables such as output, unemployment and prices and its economic effects is and will be the subject of many studies. However, even if one is not interested in the crisis as such, the pandemic has created an empirical challenge for econometricians. The extreme outliers observed in 2020 often severely distort the coefficients estimated in VARs and due to the iterative construction of impulse response functions, the latter can look quite different. For an illustration of the problem at hand, compare the IRFs of the linear VAR presented in section 5.1, with those of the same model estimated on a sample extended until Dec-2020. As it is clear from the blue lines in Figure 16, the inclusion of just 12 additional observations significantly alters the effect of a monetary policy shock, even though the baseline sample is already quite long.

Figure 16: Covid volatility: a problem and Lenza and Primiceri's solution

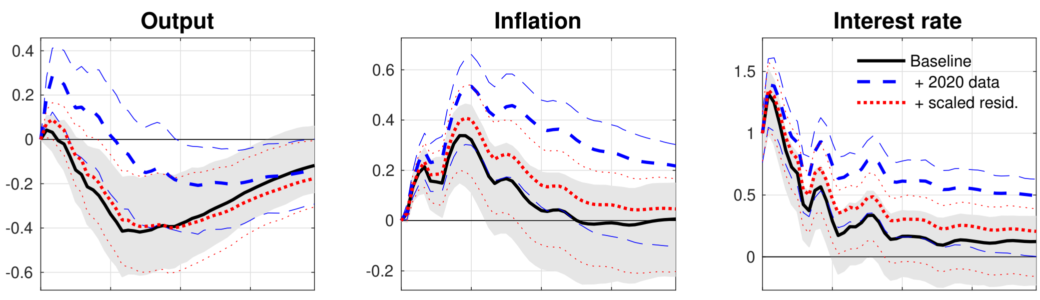

One solution is to simply exclude the respective time period. To address it more thoroughly, a "relatively simple" adjustment has been developed by Lenza and Primiceri (2020). The authors introduce the factor l_t in Equation 1 which scales the residual covariance matrix and down-weighs the extreme outliers from the pandemic:

$$Y_t = \sum_{j=1}^{p} A_j Y_{t-j} + l_t u_t. \tag{25}$$

From the start of the sample until February 2020, the period that preceded the crisis, l_t is set to 1 such that it has no effect. During the three succeeding months, where volatility in the time series was at its highest, the scaling factor assumes unique values between 0 and 1. It then gradually converges to the latter such that, eventually, it has no effect on the model. Read lines in Figure 16 show that this modification is enough to guarantee that the IRFs are quite similar to the standard linear VAR with data up until 2019.

Assuming that l_t is known, we can write 25 as a normal regression equation $Y_t = X_t\beta + l_t u_t$ where $X_t \equiv I_n \otimes x_t', x_t \equiv [Y_{t-1}', ..., Y_{t-p}']$ and $\beta \equiv vec([A_1, ..., A_p]')$. Dividing both sides of the previous equation by l_t, we get:

$$\tilde{Y}_t = \tilde{X}_t\beta + u_t, \tag{26}$$

where $\tilde{Y}_t \equiv Y_t/l_t$ and $\tilde{X}_t \equiv X_t/l_t$. Given the scaling factor l_t, $\tilde{Y}_t$ and $\tilde{X}_t$ are simple transformations of the data and thus β and Ω can be estimated ordinarily. However, in reality, the vector of the scaling factors is also unknown, meaning l_t must be estimated. This can be done with either Bayesian methods or maximum likelihood. The "extra tools" subcomponent of the toolbox uses the latter in a way that will not require any other adjustments to the codes. The estimated shocks can then be identified with any of the strategies outlined previously.[21]

9 Your project

Before we wrap up: With a 3-dimensional matrix of choices – the shock, model structure, and identification strategy – it can quickly get overwhelming. Here is an attempt at an illustration of all of the above with the most seminal papers cited for your reference.

Figure 17: Literature overview (incomplete)

	Macro shocks				
	Monetary policy	**Fiscal policy**	**Credit/finance**	**Uncertainty**	**Others**
Linear LP	Chol: Jorda 05, Haugh & Smith 12	Inst: Jorda & Taylor 15, Banerjee & Zampolli 15	Inst: Lodge & Manu 19	Chol: Biljanovska, Grigoli & Hengge 21	
Nonlinear LP	Inst: Tenreyro & Thwaites 16, Chol: Ascari & Haber 20	Inst: Ramey & Zubairy 18, Fotiou 19, Born, Müller & Pfeifer 19	Chol: Barnichon, Matthes & Ziegenbein 16	Chol: Ravenna & Cacciatore 21	Caselli & Roitman 19
Linear VAR	Chol: Christiano, Eichenbaum & Evans 99, Bernanke & Boivin 03, Coibion 12 SR: Canova & de Nicoló 02, Uhlig 05, Antolin-Diaz & Rubio-Ramirez 18 Inst: Gertler & Karadi 15 Inst+SR: Jarocinski & Karadi 20	Chol: Blanchard & Perotti 02 SR: Canova & de Nicoló 02 Inst: Ramey 12, Mertens & Ravn 14	Chol: Abbate, Eichmeier & Prieto 20, Gilchrist & Zakrajsek 12 SR: Barnett & Thomas 13 Mult.: Mumtaz, Pinter & Theodoridis 18	Chol: Bloom 09, Leduc & Liu 16 Inst: Piffer & Podstawski 18	Chol: Beaudry & Portier 16 SR: Furlanetto & Robstad 19 Inst: Känzig 21
STVAR	Chol: Iwata & Wu 06, Jackson, Owyang & Soques 18, Luigi & Huber 18, Castelnuovo & Pellegrino 18	Chol: Auerbach & Gorodnichenko 12, Caggiano, Castelnuovo, Colombo & Nodari 15	Chol: Galvao & Owyang 18		Bolboaca & Fischer 19
IVAR	Chol/SR: Aastveit, Natvik & Sola 17			Chol: Caggiano, Castelnuovo & Nodari 15, Caggiano, Castelnuovo & Pellegrino 17	
DSGE	Gali 08/15, Smets & Wouters 07, Ireland 04, Christiano, Trabandt & Walentin 11, Christiano, Eichenbaum & Trabandt 15, Gertler & Karadi 11	Leeper, Walker & Yang 13, Leeper, Traum & Walker 17, Cogan, Taylor, Wieland & Wolters 13, Fernandez-Villaverde, Guerron-Quintana, Kuester & Rubio-Ramirez 15	Kiyotaki & Moore 97, Bernanke, Gertler & Gilchrist 99, Iacoviello 05, Gilchrist, Schoenle, Sim & Zakrajsek 17 Nonlin: Guerreri & Iacoviello 17		

Task for the final paper:

- Choose a macroeconomic shock of interest and study its propagation to the economy. In doing so, you should do at least either of the two (you are welcome to do both):
  - Think about how to best identify your shock. If you have never worked with empirical time series models, it is okay to try and achieve identification with a Cholesky decomposition. Otherwise, try to go beyond it. (The goal is to learn something new.) Is there, for example, a creative instrument that can help you identify the shock? If it makes sense, compare the results of your identification scheme to the "simple" Cholesky case and discuss the strengths of your approach. Can you tell a convincing story that justifies whatever identification approach you decide to use (e.g. a particular set of sign restrictions)?

- Identify an interesting nonlinearity that could be tested in the data, estimate a nonlinear model and compare its results to the linear case. What can we learn from that exercise? Interpret your findings with economic intuition. The exercises discussed in this handbook should give you an idea of what type of nonlinearities could be interesting: Think about the state of the business cycle, constraints (on credit limits, nominal interest rates, or wage adjustment), or a nonlinearity of a particular model (e.g. difference between positive/negative, large/small shocks).

- Two approaches
  - If you have already identified an interesting topic or question, you can think about how to choose data, methodology, identification strategy, etc. to best address it. For example, you might wonder if positive and negative government spending shocks have symmetric effects? What estimation method would you need to choose in order to be able to answer that question?
  - An often promising alternative is to cross-read the literature until you stumble upon a paper you want to first replicate as precisely as possible and in a second step extend it (marginally). For example, you might find a paper that studies the effects of monetary policy on house prices using the Gertler and Karadi HFI shocks, but once you clean them for the information effect, results don't hold up. Very valuable contributions have resulted from this approach. Similarly, you could replicate a paper on data from another country and see if the results are robust.
- Two more approaches
  - The tools described here and the field more broadly allow you to estimate really complex models, facing computational challenges (e.g. the STVAR). If you like indulging in these, enjoy!
  - On the other hand, some of the best papers out there use relatively simple tools like local projections but an original idea to use the data. Both of these choices are equally valid.
- Methodology options
  - There's no need to use *all* of the described tools. Time is rather limited, so concentrate on getting one main exercise done properly instead of starting off too ambitious. There is always the possibility to do use a more complex model later on (for robustness) in case there is time.
  - The tools introduced in this handbook and provided in the toolbox of codes are some of the most widely used in empirical Macroeconomics. They are, however, not the only ones. You are free to use any of your choosing, from micro panel data to GARCH models.
- Empirical work can be cumbersome, requires careful judgement, and guarantees no success. Therefore, it's more important that your paper contains a discussion on some of the many informed choices you have to make (regarding variables included, identifying assumptions, model structure, state-dependence), rather than groundbreaking results. It might be that the results do not confirm your prior beliefs, are not in line with mainstream economic theory, or end up showing no state-dependence whatsoever. And that's okay. Describe the thought process that leads to your final product.
- Formalities
  - 36,000 characters (15 pages) for single-authored papers excl. figures and tables, double those restrictions if you work in pairs.
  - max. 8 pages of Appendices. The appendix should not contain content but only "background" information, for example a figure of the time series used, a derivation, or a robustness check. One should be able to understand all messages without reading the appendix.
  - Codes do not have to be provided.
  - Remember references and citing the most relevant literature (both in terms of methodology and the topic).
  - Work with whatever programming language you feel most comfortable. I can provide the best support in Matlab (in which most codes are still written, including mine), or R. For text processing, I recommend working with LaTeX but leave it up to your discretion.
  - The binding regulations for seminar papers and written take-home assignments can be found on KUnet

- Project description: Shortly after starting the seminar, you have to submit a project proposal of about 1 page. Timely submission is a necessary requirement for passing the seminar. The project description should contain the following elements, at least
  - outline the research question
  - motivate the shock and other variables you are going to study and why it is of interest to macroeconomists (in academia or policy).
  - outline an appropriate empirical strategy to answer the question with a special emphasis on identification and/or nonlinearities.
  - include the data you plan to use (sources, time span, frequency)

  The idea is to rigorously think through the assumptions and steps required in your analysis and less to impress the course instructor. The document is not binding. Shortly after submission, you will receive comments and suggestions to consider during your own research. However, the choices you make are yours.

Updated: August 24, 2021

Notes

[1] I would like to thank Tamás Vasi for helpful comments and suggestions, Pedro Veiga Salgado for excellent research assistance on extensions of the first version of the toolbox and Daniel Van Nguyen for careful proofreading. I am grateful to seminar participants in the spring of 2021 for their feedback, which this manuscript has greatly benefited from.

[2] Also possible: Long-run restrictions, identification via heteroskedasticity or combinations thereof

[3] Also possible: Threshold-VAR, Markov-switching VAR, Time-varying parameters (TVP-VAR), models with stochastic volatility

[4] To understand how the matrix A is used to iterate the (unidentified) VAR forward and compute the impulse response, consider first a system with 2 variables (and thus 2 shocks) and just 1 lag:

$$\begin{pmatrix} x_t \\ y_t \end{pmatrix} = \begin{pmatrix} a^{xx} & a^{xy} \\ a^{yx} & a^{yy} \end{pmatrix} \times \begin{pmatrix} x_{t-1} \\ y_{t-1} \end{pmatrix} + \begin{pmatrix} \varepsilon_t^x \\ \varepsilon_t^y \end{pmatrix}$$

Set all $x_{t-1} = y_{t-1} = \varepsilon_t^y = 0$ and shock the x variable by 1 in period 0. We thus have:

$$\begin{pmatrix} x_t \\ y_t \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$$

For all subsequent periods $h > 0$, the shocks are again set to zero, and all we need is to multiply the vector (x, y) with the dynamic matrix A:

$$\begin{pmatrix} x_{t+h} \\ y_{t+h} \end{pmatrix} = \begin{pmatrix} a^{xx} & a^{xy} \\ a^{yx} & a^{yy} \end{pmatrix} \times \begin{pmatrix} x_{t+h-1} \\ y_{t+h-1} \end{pmatrix}$$

What happens if we have $p > 1$? Every VAR(p) can be written as a VAR(1) using the so-called companion representation. Take the example of the 2-variable system with 2 lags:

$$\begin{pmatrix} x_t \\ y_t \end{pmatrix} = \begin{pmatrix} a_1^{xx} & a_1^{xy} \\ a_1^{yx} & a_1^{yy} \end{pmatrix} \times \begin{pmatrix} x_{t-1} \\ y_{t-1} \end{pmatrix} + \begin{pmatrix} a_2^{xx} & a_2^{xy} \\ a_2^{yx} & a_2^{yy} \end{pmatrix} \times \begin{pmatrix} x_{t-2} \\ y_{t-2} \end{pmatrix} + \begin{pmatrix} \varepsilon_t^x \\ \varepsilon_t^y \end{pmatrix}$$

$$\begin{pmatrix} x_t \\ y_t \\ x_{t-1} \\ y_{t-1} \end{pmatrix} = \begin{pmatrix} a_1^{xx} & a_1^{xy} & a_2^{xx} & a_2^{xy} \\ a_1^{yx} & a_1^{yy} & a_2^{xx} & a_2^{xy} \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix} \times \begin{pmatrix} x_{t-1} \\ y_{t-1} \\ x_{t-2} \\ y_{t-2} \end{pmatrix} + \begin{pmatrix} \varepsilon_t^x \\ \varepsilon_t^y \\ 0 \\ 0 \end{pmatrix}$$

which is again a VAR(1) in the vector $(x_t, y_t, x_{t-1}, y_{t-1})$ for which we can follow the above procedure. In the toolbox of codes, this procedure is performed in a function called "dynamic multipliers", whose output is the matrix C.

[5] Note here the persistently lower inflation. Because inflation is a growth rate, this implies the permanently lower price level (even once the shock has faded out) that is presented by the cited papers.

[6] However, a recent debate following Nakamura and Steinsson (2018) argues that movements within those windows have two components: An actual monetary policy component for which output falls following a tightening; but also a component labelled "central bank information effect". This is the notion that by communicating its decision, the central bank communicates new information to the market. When it raises rates, it both implements tighter policy, but also reveals that it believes the economy to

do well in the future. Jarociński and Karadi (2020) disentangle the two effects by means of sign restrictions.

[7]Most easily, we can set s^i to 1 and interpret the other impulse responses relative to it. Alternatively, we need to re-compute the standard deviation of the IV-identified shock. This is implemented in the codes.

[8]An alternative (maybe simpler) approach to the above-described instrumental variable could also be to let the series z enter as an exogenous variable to the system of equations (i.e. estimate a VARX) and study its transmission to other variables directly.

[9]A monthly series for an update of the Romer and Romer (2004) methodology is kindly provided by Silvia Miranda-Agrippino.

[10]Ramey (2016) suggests that part of this is related to the way daily data are aggregated to the monthly series in Gertler and Karadi (2015), inducing an unnecessary degree of autocorrelation.

[11]In principle, this parameter could be estimated directly, but following most of the literature, we will calibrate it to a number giving us a reasonable business cycle indicator.

[12]The code available in the toolbox to this handbook organizes the variables and settings like in the other functions, then "translates" them to format used in Auerbach and Gorodnichenko (2013) and calls the functions for the Markov chain Monte Carlo algorithm used many times in the literature. Notice that while Bayesian estimation techniques are a useful skill for a range of reasons, there are also some caveats. They are unfortunately beyond this scope of this seminar.

[13]It is possible to allow for the interacted variables to be a very simple transformation of the endogenous variables, e.g. output growth. However, this is not implemented in the code toolbox.

[14]The zero lower bound is the notion that policy-makers are reluctant to lower the nominal interest rate (far) below 0. This constraint was binding in the period following the global financial crisis (2008-2015). This period is prominently featured in the subset of histories we call recessions. The policy rate had little room to move, making it difficult to identify shocks during these quarters. Wu and Xia (2016) take yields of different financial assets along the yield curve and "translate" their movements during the zero lower bound period into a hypothetical Fed funds rate based on historical correlations. In doing so, they claim to incorporate all unconventional monetary policies that influence financial assets, including forward guidance and quantitative easing, and have a unified measure of the monetary policy stance we can interpret. Their results that unconventional monetary policy measures such as quantitative easing were as expansionary as a Fed funds rate of -3%.

[15]Therefore, C_t became y_t, products became sums, and exponents became factors. Remember the definition of growth rates as log differences: $\ln(P_{t+1}/P_t) = \ln P_{t+1} - \ln P_t = p_{t+1} - p_t = \pi_{t+1}$. We define the discount rate $\rho \equiv -\ln \beta$.

[16]Often, the system of equations needs to be log-linearized before we can simulate and estimate. To do so, we can manually calculate the steady state, perform a Taylor expansion around it and solve the log linearized system. In the NK model above, we linearize around the zero-inflation steady state. The downside of writing all variables as deviations from their natural level is that the we have no sense of movements in potential output due to technology-augmenting shocks. Alternatively, we can provide Dynare with the original equations and some information in a separate file regarding how to find the steady state.

[17]Since monetary policy is about affecting the *real* interest rate, sticky prices and therefore a flat Phillips curve make monetary policy more effective. A further note of caution: Many have argued that the Phillips curve has become flatter in recent years, by which they mean a statistical relationship between inflation and output fluctuations. This is not necessarily the same as saying that the slope in equation 20 has decreased. Consult this blog post for an excellent take at this conundrum.

[18]The reason is that the technology shock affects both the productive capacity of the economy, namely the potential level of output, and the natural rate of interest. Since our DIS equation is written in deviations from the steady state, it would affect both y and ρ. Instead, we will simply use all the data in deviations from a trend. This is not perfect, but it simplifies the points we want to make considerably. As a side effect of this, we can ignore the parameter ρ.

[19]The functions provided in the toolbox are designed to make the user understand where these hypothetical paths of variables come from. For a more rigorous discussion, you are referred to Antolin-Diaz et al. (2021) who recently developed tools to construct "structural scenarios" that can be given an economic interpretation using identified SVARs. It requires a knowledge of Bayesian econometrics which goes beyond the scope of this seminar.

[20]Conditional forecasts assume parameter stability over time but often stipulate a change in the policy reaction function, i.e. the central bank reacts differently to the economy depending on whether the ZLB is (not) binding. Put generally, changing one line of the VAR system while assuming that all other dynamic coefficients are the same can lead to misleading results (Sargent, 1979, Benati, 2010).

[21]The code and further explanations on the procedure are publicly available from Lenza and Primiceri (2020).

References

Abbate, Angela, Sandra Eickmeier, and Esteban Prieto, "Financial Shocks and Inflation Dynamics," *Swiss National Bank Working Paper*, 2020.

Antolin-Diaz, Juan and Juan F. Rubio-Ramirez, "Narrative Sign Restrictions," *American Economic Review*, 2018, *108*, 2802–2829.

__ , Ivan Petrella, and Juan F. Rubio-Ramirez, "Structural Scenario Analysis with SVARs," *Journal of Monetary Economics*, 2021, *117*, 798–815.

Ascari, Guido and Timo Haber, "Non-Linearities, State-Dependent Prices and the Transmission Mechanism of Monetary Policy," *The Economic Journal*, forthcoming.

Auerbach, Alan and Yuriy Gorodnichenko, "Measuring the Output Responses to Fiscal Policy," *American Economic Journal: Economic Policy*, 2012, *4* (2), 1.

__ and __ , "Corrigendum: Measuring the Output Responses to Fiscal Policy," *American Economic Journal: Economic Policy*, 2013, *5* (3), 320–322.

Auerbach, Alan J. and Yuriy Gorodnichenko, "Fiscal Multipliers In Recession and Expansion," in "Fiscal Policy After the Financial Crisis," National Bureau of Economic Research, Inc., 2012, pp. 63–98.

Beaudry, Paul and Franck Portier, "Stock Prices, News, and Economic Fluctuations," *American Economic Review*, 2006, *96* (4), 1293–1307.

Benati, Luca, "Are Policy Counterfactuals Based on Structural VAR's Reliable?," *European Central Bank Working Paper Series 1188*, 2010.

Berger, David, Konstantin Milbradt, Fabrice Tourre, and Joseph Vavra, "Mortgage Prepayment and Path-Dependent Effects of Monetary Policy," *American Economic Review*, forthcoming.

Bernanke, Ben and Jean Boivin, "Monetary Policy in a Data-Rich Environment," *Journal of Monetary Economics*, 2003, *50*, 5.

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

Caggiano, Giovanni, Efrem Castelnuovo, and Giovanni Pellegrino, "Estimating the Real Effects of Uncertainty Shocks at the Zero Lower Bound," *European Economic Review*, 2017, *100*, 257–272.

__ , __ , Valentina Colombo, and Gabriela Nodari, "Estimating Fiscal Multipliers: News from a Nonlinear World," *Economic Journal*, 2015, *125* (584), 746–776.

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

Caselli, Francesco and Agustin Roitman, "Nonlinear Exchange-rate Pass-through in Emerging Markets," *International Finance*, 2019, *22* (3), 279–306.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans, "Monetary Policy Shocks: What Have We Learned and to What End?," in J. B. Taylor and M. Woodford, eds., *Handbook of Macroeconomics*, 1999, pp. 65–148.

Dedola, Luca, Georgios Georgiadis, Johannes Gräb, and Arnaud Mehl, "Does a Big Bazooka Matter? Quantitative Easing Policies and Exchange Rates," *Journal of Monetary Economics*, 2021, *117*, 489–506.

Fry, Renee and Adrian Pagan, "Sign Restrictions in Structural Vector Autoregressions: A Critical Review," *Journal of Economic Literature*, 2011, *49* (4), 938–960.

Furlanetto, Francesco and Ørjan Robstad, "Immigration and the Macroeconomy: Some New Empirical Evidence," *Review of Economic Dynamics*, 2019, *34*, 1–19.

Galí, Jordi, "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?," *American Economic Review*, 1999, *89* (1), 249–271.

__ , *Monetary Policy, Inflation, and the Business Cycle*, Princeton University Press, 2015.

Gertler, Mark and Peter Karadi, "Monetary Policy Surprises, Credit Costs, and Economic Activity," *American Economic Journal: Macroeconomics*, 2015, *7* (1), 44–76.

Gilchrist, Simon and Egon Zakrajšek, "Credit Spreads and Business Cycle Fluctuations," *American Economic Review*, 2012, *102* (4), 1692–1720.

Guerrieri, Luca and Matteo Iacoviello, "OccBin: A Toolkit for Solving Dynamic Models With Occasionally Binding Constraints Easily," *Journal of Monetary Economics*, 2015, *70*, 22–38.

Holm, Martin Blomhoff, Pascal Paul, and Andreas Tischbirek, "The Transmission of Monetary Policy under the Microscope," *Journal of Political Economy*, forthcoming.

Iacoviello, Matteo, "House Prices, Borrowing Constraints and Monetary Policy in the Business Cycle," *American Economic Review*, 2005, *95* (3), 739–764.

Iwata, Shigeru and Shu Wu, "Estimating Monetary Policy Effects when Interest Rates are Close to Zero," *Journal of Monetary Economics*, 2006, *53* (7), 1395–1408.

Jarociński, Marek and Peter Karadi, "Deconstructing Monetary Policy Surprises: The Role of Information Shocks," *American Economic Journal: Macroeconomics*, 2020, *12* (2), 1–43.

Jordà, Oscar, "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, 2005, *95* (1), 161–182.

Känzig, Diego, "The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements," *American Economic Review*, 2021, *111* (4), 1092–1125.

Kilian, Lutz and Helmut Lütkepohl, *Structural Vector Autoregressive Analysis*, Cambridge University Press, 2017.

Koop, Gary, M Hashem Pesaran, and Simon M Potter, "Impulse Response Analysis in Nonlinear Multivariate Models," *Journal of Econometrics*, 1996, *74* (1), 119–147.

Leduc, Sylvain and Zheng Liu, "Uncertainty Shocks are Aggregate Demand Shocks," *Journal of Monetary Economics*, 2016, *82* (C), 20–35.

Lenza, Michele and Giorgio E Primiceri, "How to Estimate a VAR after March 2020," *NBER Working Paper*, 2020.

Lhuissier, Stéphane and Benoît Nguyen, "The Dynamic Effects of the ECB's Asset Purchases: A Survey-Based Identification," *Banque de France Working Paper*, 2021.

Luigi, Clara De and Florian Huber, "Debt Regimes and the Effectiveness of Monetary Policy," *Journal of Economic Dynamics and Control*, 2018, *93*, 218–238.

Lütkepohl, Helmut and Aleksei Netšunajev, "Structural Vector Autoregressions with Smooth Transition in Variances: The Interaction between U.S. Monetary Policy and the Stock Market," 2014. DIW working paper.

Mertens, Karel and Morten O. Ravn, "Empirical Evidence on the Aggregate Effects of Anticipated and Unanticipated US Tax Policy Shocks," *American Economic Journal: Economic Policy*, 2014, *4* (2), 145–181.

Miranda-Agrippino, Silvia, "Unsurprising Shocks: Information, Premia, and the Monetary Transmission," *Bank of England Working Papers*, 2016, *626.*

Nakamura, Emi and Jón Steinsson, "High Frequency Identification of Monetary Non-Neutrality: The Information Effect," *Quarterly Journal of Economics*, 2018, *133* (3), 1283–1330.

Olivei, Giovanni P and Silvana Tenreyro, "The Timing of Monetary Policy Shocks," *American Economic Review*, 2007, *97* (3), 636–663.

Ottonello, Pablo and Thomas Winberry, "Financial Heterogeneity and the Investment Channel of Monetary Policy," *Econometrica*, 2020, *88* (6), 2473–2502.

Piffer, Michele and Maximilian Podstawski, "Identifying Uncertainty Shocks Using the Price of Gold," *The Economic Journal*, 2017, *128* (616), 3266–3284.

Plagborg-Møller, Mikkel and Christian K Wolf, "Local Projections and VARs Estimate the Same Impulse Responses," *Econometrica*, 2021, *89* (2), 955–980.

Ramey, Valerie, "Macroeconomic Shocks and Their Propagation," in J. B. Taylor and H. Uhlig, eds., *Handbook of Macroeconomics*, 2016, pp. 71–162.

_ and Sarah Zubairy, "Government Spending Multipliers in Good Times and in Bad," *Journal of Political Economy*, 2018, *126* (2), 850–901.

Romer, Christina and David Romer, "A New Measure of Monetary Shocks: Derivation and Implications," *American Economic Review*, 2004, *94* (4), 1055–1084.

Sargent, Thomas J., "Estimating Vector Autoregressions Using Methods Not Based on Explicit Economic Theories," *Federal Reserve Bank of Minneapolis Quarterly Review*, 1979, pp. 8–15.

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

Sims, Eric, Jing Cynthia Wu, and Ji Zhang, "The Four Equation New Keynesian Model," *The Review of Economics and Statistics*, forthcoming.

Smets, Frank and Rafael Wouters, "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, 2007, *97* (3), 586–606.

Tenreyro, Silvana and Gregory Thwaites, "Pushing on a String: US Monetary Policy Is Less Powerful in Recessions," *American Economic Journal: Macroeconomics*, 2016, *8* (4), 43–74.

Uhlig, Harald, "What are the Effects of Monetary Policy? Results from an Agnostic Identification Procedure.," *Journal of Monetary Economics*, 2005, *52*, 381–419.

Wu, Jing Cynthia and Fan Dora Xia, "Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound," *Journal of Money, Credit, and Banking*, 2016, *48* (2-3), 253–291.