PhD Macroeconomics 1
6. More on VARs and Alternatives

Karl Whelan

School of Economics, UCD

Autumn 2023

Roadmap

1. Monetary policy VARs
2. Fiscal policy VARs
3. VARs with long-run restrictions
4. An alternative to VARs: The local projection method
5. An alternative to VARs: The narrative method

Part I

Monetary Policy VARs

Using VARs to Analyse Monetary Policy

- One of the key issues in macroeconomics is how monetary policy affects the economy.
- Central banks would ideally like to have a detailed picture of exactly how their decisions affect the economy: If we raise interest rates by $x + \epsilon$ instead of x what would be the impact on GDP and inflation in one quarter, two quarters, three quarters ... eight quarters?
- But figuring out the effects of monetary policy requires dealing with reverse causality issues: Central banks influence the economy but the actions of central banks also respond to what is going on in the economy.
- VARs seem to provide an ideal way to answer these questions.
    - They provide impulse responses that tell you how policy actions impact variables over time.
    - The Cholesky decomposition method provides a way to model simultaneous causality between variables.
- Let's take a look at how we would use a VAR to examine the impacts of monetary policy.

A Monetary Policy VAR

- Let's go back to the VAR we showed at the end of the previous lecture notes with 3 variables.
  1. Inflation, defined as the annualised percentage change in the deflator for personal consumption expenditures.
  2. Unemployment, defined as the unemployment rate in the monthly household survey.
  3. The average federal funds rate for the month. This is the Fed's "policy rate".
- Stock and Watson's 2001 *Journal of Economic Perspectives* paper is both a useful introduction to VAR methods and also provides a substantive application. They estimate a VAR with these variables using quarterly data.
- We are going to use our longer monthly dataset, taken from FRED-MD, to re-do their analysis.

Identification Strategy

- We have monthly data on inflation (π_t), the unemployment rate (u_t) and the federal funds rate (i_t).
- We follow Stock and Watson in assuming a "lower-triangular" causal chain of contemporaneous interactions of the form

$$AZ_t = \begin{pmatrix} a_{11} & 0 & 0 \\ a_{21} & a_{22} & 0 \\ a_{31} & a_{32} & a_{33} \end{pmatrix} \begin{pmatrix} \pi_t \\ u_t \\ i_t \end{pmatrix} = BZ_{t-1} + \epsilon_t$$

- The identifying assumptions
  1. Inflation depends only on lagged values of the other variables (perhaps motivated by the idea of sticky prices.)
  2. Unemployment depends on contemporaneous inflation but not the funds rate.
  3. The funds rate depends on both contemporaneous inflation and unemployment. (The Fed using its knowledge about the current state of the economy when it is setting interest rates).

Impulse Response Graphs for the Monetary Policy VAR

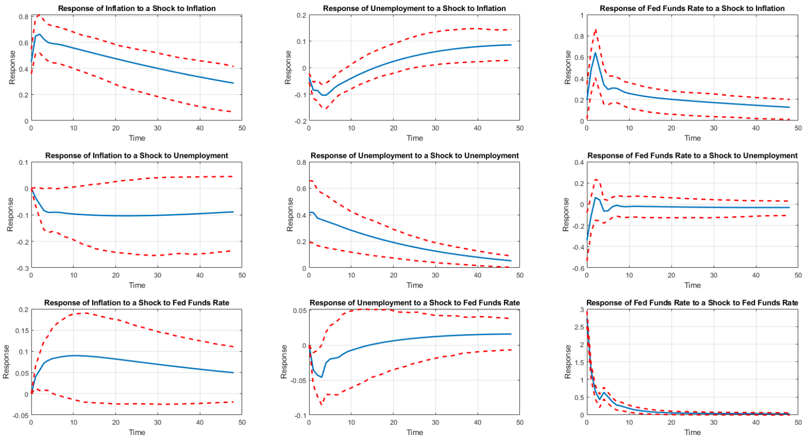

A Puzzle?

- Some of these results make sense. The fed funds rate rises with a positive shock to inflation and falls with a positive shock to unemployment. A positive shock to unemployment reduces inflation and a positive interest rate shock eventually raises the unemployment rate.
- However, the response of the inflation rate to a shock to the Fed funds rate is puzzling: The interest rate increase seems to raise the inflation rate.
- This result (labelled the "price puzzle" in line with economists' gift for giving things bad names) has been obtained in many VAR studies. It provides a good illustration of the potential limitations of VAR analysis.
- One explanation is that Fed is acting on information not captured in the VAR (for example, information about commodity prices) and that this information may provide signals of future inflationary pressures.
- So interest rate increases may occur just before an increase in inflation. The VAR may be capturing this pattern and confusing causation and correlation.
- Adding commodity prices to the VAR has sometimes been found to eliminate the "price puzzle." But is this improving the specification or is it data mining to find the result we want?

A Second Identification

- The results from the Stock-Watson identification generally make sense but you could imagine other possible identifications.
- The next slide shows results from an alternative ordering with interest rates first, unemployment second, and inflation last.
- You could argue for this case on the grounds that the Fed can only respond to the economy with a lag because it takes a while to receive data about the current state of the economy (so they are reacting to lagged information) but that inflation should be able to respond immediately to economic events.
- This may sound reasonable enough but the results from this identification make even less sense.
- Which identification a researcher settles on may depend on how "sensible" they believe results are. This may be problematic.

IRFs From Recursive VAR, Second Identification

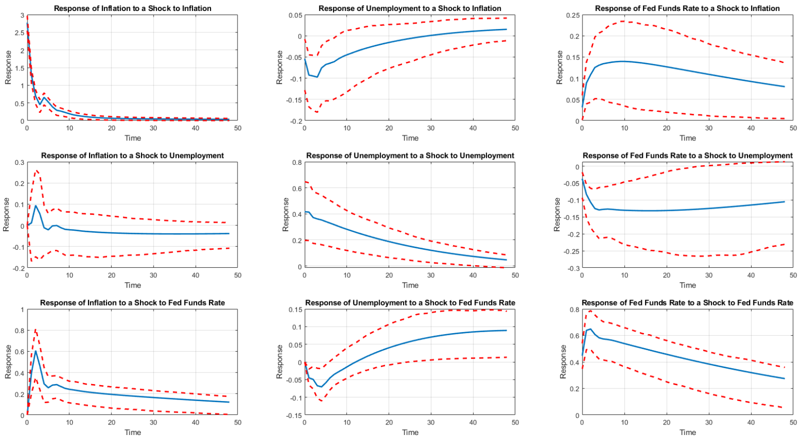

Rudebusch's Critique

- A lot of interesting material on monetary policy VARs can be found in a 1998 exchange between Glenn Rudebusch and Chris Sims.
- Rudebusch's paper contains a strong critique of VARs used to assess the effects of monetary policy. Among his points:
  1. The VARs ignore changes over time in the formulation of monetary policy.
  2. They use final published data instead of the preliminary estimates the Fed has available when it makes decisions.
  3. They greatly underestimate the information available to the Fed when it takes decisions.
  4. They incorporate long lags but he argues it is not credible that the Fed responds to information from over a year prior to taking a decision.
  5. The monetary policy shocks don't look anything like the surprise element of monetary policy decisions obtained from looking at financial contracts like fed funds futures.
  6. Models with very different monetary policy shocks report similar IRFs, suggesting that perhaps the models have been data-mined to give these answers.

The Rudebusch-Sims Debate

- Sims responded in detail to Rudebusch's critiques. Some of his points were as follows.
  1. VAR models may differ in their shocks but agree on their effects. For example, one model may include more variables in a supply equation than another so its supply shocks are, by construction, smaller in size but both models could still capture roughly the same effect of a shock to supply.
  2. Financial market "surprises" are not necessarily the best measures of the exogenous element of monetary policy. A Fed governor could give a speech the week before an FOMC meeting indicating that the Fed is going to raise rates (even if this isn't predicted by inflation or GDP or other standard VAR variables.) When this rate increase happens, we'd like to know what effect it has even though, on the day, it is not a surprise for financial markets.
  3. Issues like "time invariance, linearity, and variable selection are universal in macroeconomic modeling" and are not special to VARs.
- And Rudebusch responded back ...
- The exchange is well worth reading in full as it sheds light on a lot of the issues that economists need to think about when doing VAR analysis.

Cochrane's (2023) Discussion

- In a 2023 blogpost, John Cochrane, a leading macroeconomist has a critique of using VARs to assess monetary policy.
- Cochrane makes many points but I will note two here.
  - **What are the shocks?**: "*What's a "shock" anyway? The concept is that the Fed considers its forecast of inflation, output and other variables it is trying to control, gauges the usual and appropriate response, and then adds 25 or 50 basis points, at random, just for the heck of it. The question VARs try to answer is the same: What happens to the economy if the Fed raises interest rates unexpectedly, for no particular reason at all? But the Fed never does this. Ask them. Read the minutes. The Fed does not roll dice. They always raise or lower interest rates for a reason, that reason is always a response to something going on in the economy, and most of the time how it affects forecasts of inflation and employment.*
  - **Lack of robustness**: He notes that the impact of interest rates on inflation varies very widely depending on the specifications chosen. It is hard to find trustworthy numbers from VAR analysis.

Summary

- The identification issues for using VARs to analyse the effects of monetary policy on the economy are complex.
- Central banks make decisions based on what is happening in the economy and figuring out "causal orderings" involving these decisions is not easy.
- It can also be hard to distinguish "pure shocks"—deviations from how central banks usually respond to the economy—from changes in the central bank's reaction function.
- Ultimately, it is hard to use a mechanical method like a VAR to highlight random "exogenous" changes in monetary policy.
- So we probably should not be too surprised that the results from monetary policy VARs are unsatisfactory and generally not robust.

Part II

Fiscal Policy VARs

A Fiscal Policy VAR

- Blanchard and Perotti (2002) examined a three-variable VAR using quarterly US data from 1947 to 1997: Federal tax revenues, federal government spending and GDP. Their VAR can be written in reduced form as T_t

$$Z_t = \begin{pmatrix} T_t \\ G_t \\ X_t \end{pmatrix} = AZ_{t-1} + e_t \qquad e_t = \begin{pmatrix} e_t^T \\ e_t^G \\ e_t^X \end{pmatrix}$$

  where T_t, G_t and X_t are real per capita tax revenue, government spending and GDP and e_t contains the reduced-form shocks.

- They assume the reduced-form shocks are related to the structural shocks ϵ_t^T, ϵ_t^G, ϵ_t^X as follows

$$\begin{aligned} e_t^T &= a_1 e_t^X + a_2 e_t^G + \epsilon_t^T \\ e_t^G &= b_1 e_t^X + b_2 e_t^T + \epsilon_t^G \\ e_t^X &= c_1 e_t^T + c_2 e_t^G + \epsilon_t^X \end{aligned}$$

  where ϵ_t^T, ϵ_t^G and ϵ_t^X are uncorrelated structural shocks.

Identification Strategy

- Blanchard and Perotti do not use the Cholesky decomposition. Instead, they use prior information to set three different contemporaneous effect coefficients:
  1. They use separate information on tax elasticities to set a specific positive value for the contemporaneous effect of GDP on tax revenues, a_1. It varies over time, averaging $a_1 = 2.08$.
  2. They assume no within-quarter effect of GDP on government spending so they set $b_1 = 0$. They note "*Direct evidence on the conduct of fiscal policy suggests that it takes policymakers ... more than a quarter to learn about a GDP shock, decide what fiscal measures, if any, to take in response, pass these measures through the legislature, and actually implement them.*"
  3. They then swap between setting the contemporanous effect of taxes on spending equal to zero ($b_2 = 0$) and setting the contemporanous effect of spending on taxes equal to zero ($a_2 = 0$) and report that results are similar when either is used.
- BP assume that taxes and spending can both affect GDP within the same quarter, i.e. they place GDP last in the ordering as we describe it.

Results

- Blanchard and Perotti presented two different sets of results, differing in whether they treat their variables as having a deterministic trend or a stochastic trend (the latter was probably more reliable).
- Both sets of results provide some support for Keynesian positions on fiscal policy. They found that tax increase shocks have a negative effect on GDP and also that government spending increase shocks have a positive effect on GDP (though these results were not statistically significant).
- Since 2002, there has been lots of empirical macro research on fiscal policy.
- Valerie Ramey's 2019 *Journal of Economic Perspectives* article provides a summary of this research. She starts by updating Blanchard and Perotti's VAR and reporting statistically significant positive impacts of government spending shocks on GDP (see the graph a few pages down).
- She reports "*On average, government purchases multipliers are likely to be between 0.6 and 1. Narrative-based time series estimates point to tax rate change multipliers between -2 and -3.*"
- Empirical research on fiscal policy seems to have been more successful than for monetary policy.

Impluse Responses to Higher Taxes

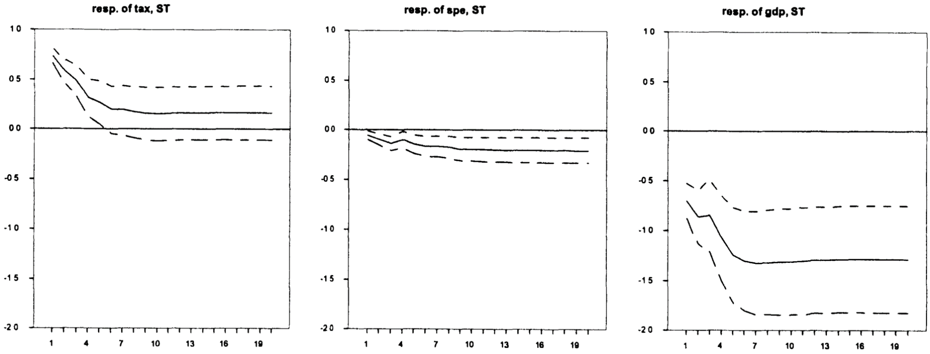

Impluse Responses to Higher Government Spending

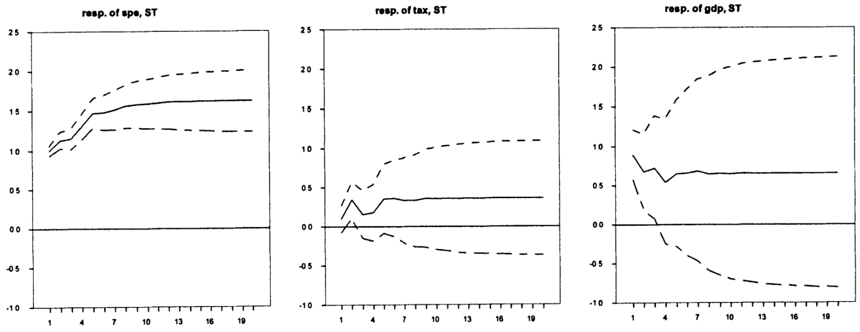

Ramey's Updated Estimate of Government Spending IRF

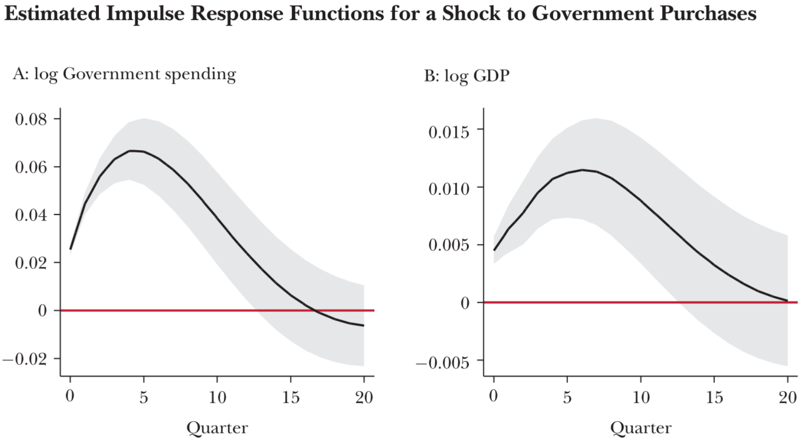

Part III

VARs With Long-Run Restrictions

An Alternative Approach: Long-Run Restrictions

- The identifying assumptions in the recursive VAR approach require knowledge of how certain variables react in an instantaneous way to certain shocks.
- Sometimes, because certain variables are "sluggish" or because information about some variables is only available with a lag, we can be pretty confident about these restrictions. But often they are pure guesswork.
- And economic theory gives very little guidance.
- In fact, economic theory usually tells us a lot more about what will happen in the longer-run, rather than exactly what will happen today.
- For instance, theory tells us that whatever positive aggregate demand shocks do in the short-run, in the long-run they have no effect on output and a positive effect on the price level.
- This suggests an alternative approach: Use these theoretically-inspired long-run restrictions to identify shocks and impulse responses.
- I will explain one of these methods over the next four slides. The general idea is more important than the technical details. (I won't ask for this derivation in the final exam).

Information in the Reduced-Form Covariance Matrix

- Consider the VAR model

$$Z_t = BZ_{t-1} + C\epsilon_t$$

where the covariance matrix of the structural shocks is

$$E(\epsilon_t \epsilon_t') = \begin{pmatrix} E(\epsilon_1^2) & E(\epsilon_1 \epsilon_2) \\ E(\epsilon_1 \epsilon_2) & E(\epsilon_2^2) \end{pmatrix} = I$$

so the structural shocks are uncorrelated and have unit variance (this is just a harmless normalization).

- Note that the covariance matrix of the observed reduced-form errors is

$$\Sigma = E(e_t e_t') = E\{(C\epsilon_t)(C\epsilon_t)'\} = CE(\epsilon_t \epsilon_t')C' = CC'$$

- Thus, the observed covariance structure of the reduced-form shocks tells us something about how they are related to the uncorrelated, unit-variance, structural shocks.

Calculating Long-Run Effects in an SVAR

- Suppose $Z_t = (\Delta y_t, \Delta x_t)'$
- Then the long-run effect of the shock on y_t is the sum of its effects of Δy_t, Δy_{t+1}, Δy_{t+2} and so on.
- The long-run effect is the sum of the impulse responses.
- The impulse responses for the model $Z_t = BZ_{t-1} + C\epsilon_t$ are
  1. C in impact period.
  2. BC after one period.
  3. B^2C after two periods, ... B^nC after n periods.
- Long-run level effects are $D = \left(I + B + B^2 + B^3 + ...\right) C$.
- If eigenvalues of B are inside unit circle then $I + B + B^2 + B^3 + ... = (I - B)^{-1}$.
- This is the matrix equivalent of the multiplier formula $1 + c + c^2 + c^3 + ... = \frac{1}{1-c}$.
- So the long-run responses are $D = (I - B)^{-1} C$.

The Blanchard-Quah Method: I

- This method was introduced by Blanchard and Quah (1989).
- Now note that $DD' = (I - B)^{-1} CC' \left((I - B)^{-1} \right)'$
- But two slides ago, we established that $CC' = \Sigma$, the covariance matrix of the reduced-form shocks, which can be estimated.
- So $DD' = (I - B)^{-1} \Sigma \left((I - B)^{-1} \right)'$ and this matrix can also be calculated because we can estimate B from OLS estimation of the reduced-form VAR.
- Now make a restriction about the long-run effects described in D: Assume that D is lower triangular: Only the first shock has a long-run effect on the first variable, and only the first and second shocks have long-run effects on the second variable and so on.
- In the two variable case, this is just

$$D = \begin{pmatrix} d_{11} & 0 \\ d_{21} & d_{22} \end{pmatrix}$$

The Blanchard-Quah Method: II

- $DD' = (I - B)^{-1} \Sigma \left((I - B)^{-1} \right)'$ is a symmetric matrix (the i, j entry is identical to the j, i entry)
- All symmetric matrices have a unique lower-diagonal matrix D so that DD' equals the symmetric matrix. This is known as the **Cholesky factor** of the symmetric matrix.
- D can be calculated directly using various software as the Cholesky factor of the known matrix $(I - B)^{-1} \Sigma \left((I - B)^{-1} \right)'$.
- Now remember that $D = (I - B)^{-1} C$.
- So the crucial matrix C defining the structural shocks can then be calculated as

$$C = (I - B)D$$

- Now, we can calculate the impulse response functions to the structural shocks.

Blanchard-Quah: Identifying Supply and Demand Shocks

- Blanchard-Quah (1989) used a two-variable VAR in the log-difference in GDP Δy_t and the unemployment rate U_t.
- Unemployment was entered in levels form. Because the VAR is estimated to be stationary (eigenvalues inside unit circle) both structural shocks have zero long-run effect on the unemployment rate.
- The lower-diagonal assumption thus meant that of the two structural shocks only one of them could have a long-run effect on the level of output. BQ labelled this the "supply" shock while the shock that has no effect on long-run output was labelled the "demand" shock.
- The relative importance of supply versus demand shocks in determining output is a long-running theme in macroeconomics. Keynesians emphasize the importance of demand shocks while more classically-oriented economists, such as advocates of the real business cycle approach, see supply shocks as being more important.
- BQ's results implied that demand shocks were responsible for the vast majority of short-run fluctuations.

Galí (1999): Technology Shocks and Hours Worked

- BQ's formulation is restrictive: The assumption that neither supply or demand shocks can change the unemployment rate in the long-run may not be correct.
- Galí's paper applied a similar analysis to BQ, but for a formulation that moved a bit closer to the debate about real business cycle models and their predictions for the labour market.
- RBC models assume technology shocks drive the business cycle and explain why hours worked are higher in booms than in recessions: Better to work when you are productive than unproductive.
- Galí's VAR featured the log-difference of output per hour worked (labour productivity), Δz_t and the log-difference of hours worked, Δn_t.
- The lower-diagonal assumption about long-run responses now means that the supply shock (now called the "technology" shock) can affect productivity in the long-run, while the non-techology shock cannot.
- The model lets the data dictate the long-run effects of technology and non-technology shocks on hours worked.
- An updated version of Galí's IRFs is on the next page from a paper I published in *Journal of Macroeconomics*.

Replication of Galí's Results (with 10% and 90% Bootstrapped Error Bands)

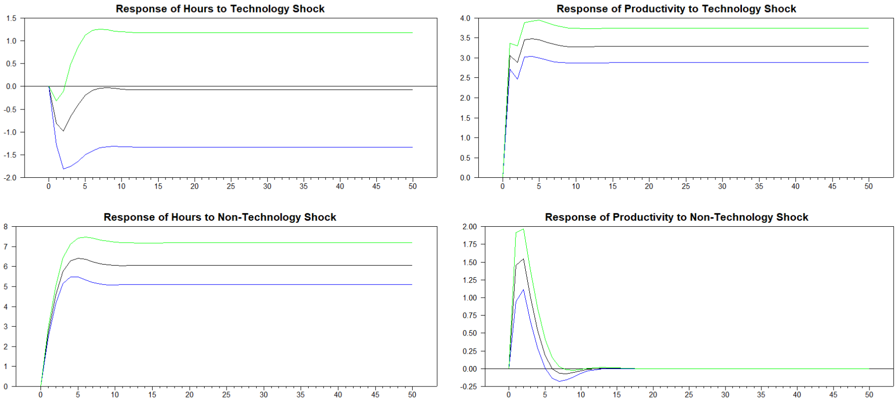

Interpreting Galí's Results

1. Positive non-technology shocks cause both output and productivity to rise in the short-run.
   - This provides evidence that short-run cyclical movements in productivity are not just due to technology shocks.
   - One explanation is costs of adjusting labour input. Rather than hire new labour in a boom, firms may existing workers temporarily work a bit harder. In recessions, employed labour is more likely to be under-utilized.
2. Positive non-technology shocks raise labour input.
   - The interpretation of this result unclear. Perhaps positive demand shocks that encourage workers to join the labour force make people more likely to stay employed.
3. Technology shocks cause productivity to go up but hours to go down.
   - Interpretation: More efficiency means a given level of output can be supplied with less labour. It takes time for the technology shock to fully pass through to output.
   - Bad news for technology-driven stories of labour market fluctuations such as real business cycle models.

Part IV

Alternatives: The Local Projection Method and the Narrative Approach

The Local Projection Method

- Introduced by Oscar Jorda (2005), the local projection method is a more flexible way of estimating impulse responses than VARs.
- When we estimate a VAR of the form $Z_t = AZ_{t-1} + \epsilon_t$, the impulse responses are found by calculating A, A^2, A^3 and so on.
- All IRFs depend on the estimated dynamic relationships between Z_t and Z_{t-1}.
- However, this may miss various complexities and non-linearities such that the relationship between, for example, Z_{t+5} and Z_{t-1} might not be correctly estimated by calculating A from a first-order VAR and then calculating A^5.
- Jorda recommends estimating h different impulse responses for different horizons by applying h different regressions of the form

$$Z_{t+h} = BZ_{t-1} + \epsilon_{t-h}$$

  so the $(h+1)$-step ahead impulse responses are calculated as $\hat{B}$. The regression output will automatically provide the standard errors for the impulse responses.
- This can be extended to have additional lagged variables in the regression but the $(h+1)$-step ahead impulse responses would still be B.

The Narrative Approach to Monetary Policy

- There are other ways to try to get at "truly exogenous" movements in monetary policy.
- For example, Christina Romer and David Romer have a written a series of papers outlining a "narrative approach," most recently their 2023 *AER* paper.
- They (2023) look at each FOMC meeting and attempt to get at a measure of monetary shocks that "*should be relatively free of both endogenous and anticipatory actions.*"
- Their current list of shocks is on the next page. There are no shocks after 1988 but they believe 2022 will eventually be counted as a shock.
- They use the local projection method to estimate the effects of a shock dummy variable S_t on outcome variables Y_t

$$Y_{t+h} = \alpha^h + \beta^h S_t + \sum_{k=1}^{K} \psi_k^h S_{t-k} + \sum_{k=1}^{K} \theta_k^h Y_{t-k} + \epsilon_{t+h}$$

where the sequence of the estimated β^h values for the various horizons is an estimate of the response of the outcome variable to a realization of 1 for the dummy variable

Romer and Romer Monetary Policy Shocks

TABLE 2—MONETARY POLICY SHOCKS, 1946–2016

New dates			Original dates		
October	1947	(–)	October	1947	(–)
August	1955	(–)	September	1955	(–)
September	1958	(–)			
December	1968	(–)	December	1968	(–)
January	1972	(+)			
April	1974	(–)	April	1974	(–)
August	1978	(–)	August	1978	(–)
October	1979	(–)	October	1979	(–)
May	1981	(–)			
December	1988	(–)	December	1988	(–)

Notes: Contractionary shocks are denoted (–) and expansionary shocks are denoted (+). In setting our original dates, we did not have a classification for expansionary shocks.

Romer and Romer (2023): Effect of a contractionary monetary policy shock on unemployment

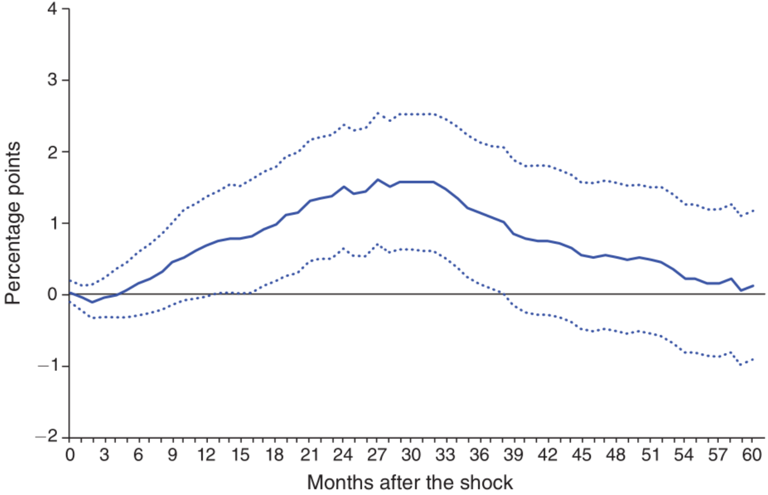

Romer and Romer (2023): Effect of a contractionary monetary policy shock on real GDP

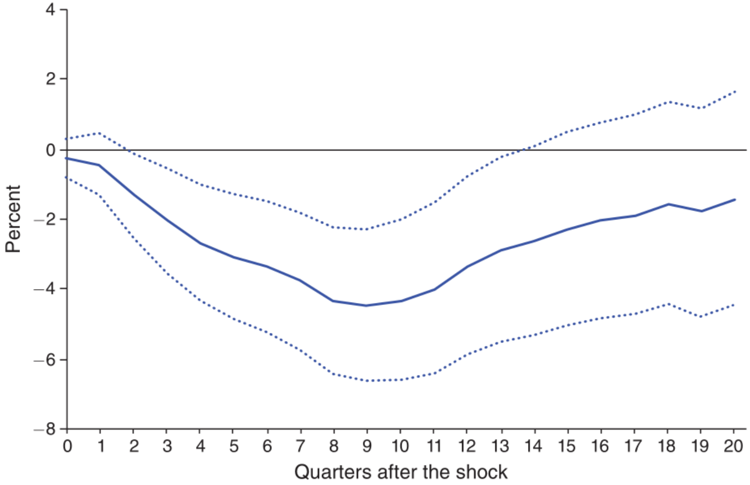

Romer and Romer (2023): Effect of a contractionary monetary policy shock on inflation

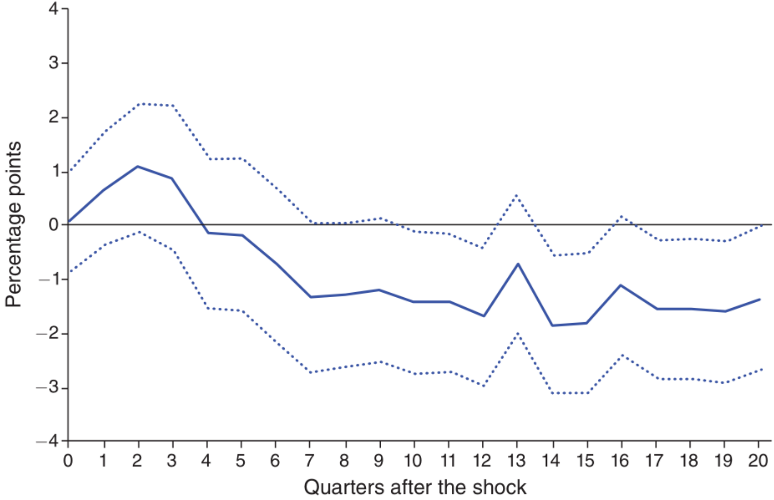

The Narrative Approach to Fiscal Policy

- Romer and Romer (2010) also apply their narrative approach to looking for exogenous changes in tax policy.
- They describe their approach as follows: "*There exists a vast narrative record describing the history and motivation of tax policy changes. We first use this narrative history to separate legislated tax changes from those arising from nonpolicy developments. We then use the information on motivation to separate the legislated tax changes into those that are likely to be contaminated by other developments affecting output, and those that can legitimately be used to measure the macroeconomic effects of tax changes.*
- They then assess the impact of these exogenous shocks on the economy and report relatively large and persistent contractionary effects of tax increases on GDP.

Romer and Romer Tax Policy Shocks

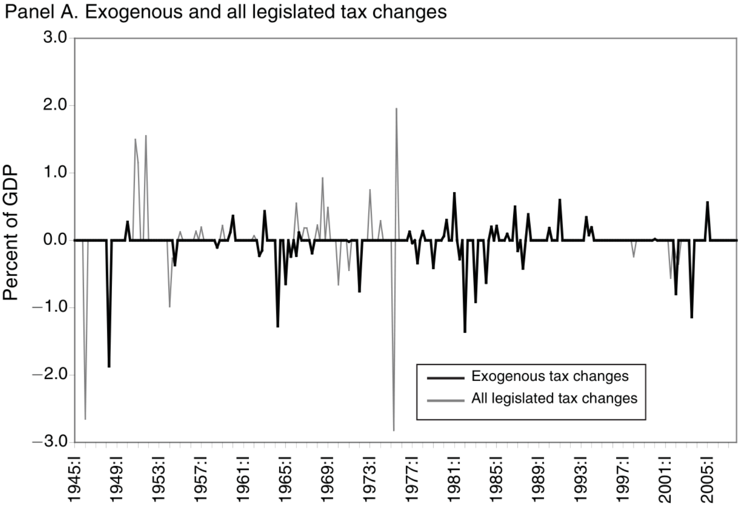

Romer and Romer (2023): Effect on GDP of a contractionary tax increse of 1% of GDP

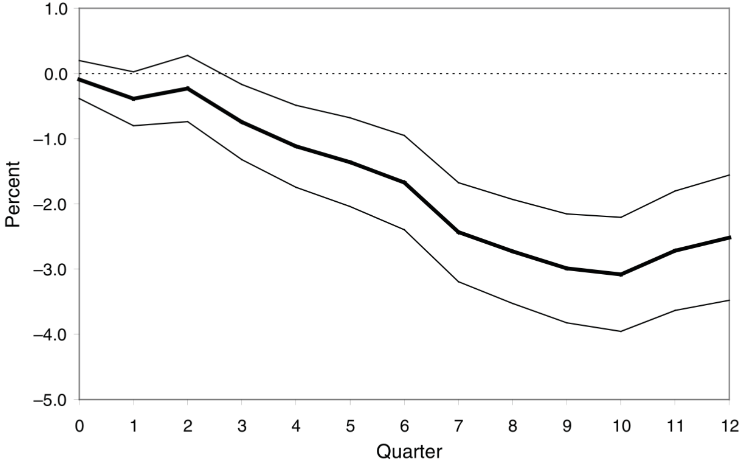

FIGURE 4. ESTIMATED IMPACT OF AN EXOGENOUS TAX INCREASE OF 1 PERCENT OF GDP ON GDP