Identifying Aggregate Demand and Supply Shocks Using Sign Restrictions and Higher-Order Moments

Geert Bekaert*

Columbia University and CEPR

Eric Engstrom†

Board of Governors of the Federal Reserve System

Andrey Ermolov‡

Fordham University

October 2, 2022

*gb241@columbia.edu

†eric.c.engstrom@frb.gov. The views expressed herein do not necessarily represent those of the Federal Reserve System, its Board of Governors, or staff.

‡aermolov1@fordham.edu

Abstract

We use information in higher-order moments to identify aggregate supply and aggregate demand shocks for the U.S. economy. Traditional methods based on sign restrictions and/or second-order moments yield only "set" or "interval" identification but higher-order moments are shown to considerably aid identification. Aggregate supply shocks dominated recessions in the 1970s and early 1980s, while aggregate demand shocks dominated most later recessions. The Great Recession of 2008-2009 and the pandemic-induced recession of 2020 exhibited both large negative aggregate demand and large negative aggregate supply shocks.

JEL Classification: E31, E32, E43, E44

Keywords: business cycles, Great Moderation, AS/AD shocks, kurtosis, Covid pandemic

I Introduction

Distinguishing aggregate supply (AS) shocks from aggregate demand (AD) shocks has long been a central topic in empirical macroeconomics (e.g., Shapiro and Watson, 1988, Blanchard and Quah, 1989), in part because adverse demand shocks require different monetary and fiscal policy responses than do adverse supply shocks. Accommodative monetary policy may help offset a shortfall of aggregate demand but may be impotent or counterproductive in the face of a negative aggregate supply shock. Recently, the Covid-19 pandemic and its aftermath have brought the important differential effects of AS and AD shocks back to the fore of economic and political debate. While economic lockdowns associated with the pandemic arguably represented large negative demand shocks, impaired supply chains certainly affected aggregate supply, and various fiscal responses likely affected both aggregate supply and aggregate demand.

We propose a new identification methodology to distinguish AS and AD shocks. We only use stylized Keynesian definitions of aggregate supply and demand shocks, as motivated, for instance by Blanchard (1989) and Gali (1992). In particular, we define aggregate supply shocks as shocks that move inflation and real activity in opposite directions; conversely, demand shocks are defined as innovations that move inflation and real activity in the same direction. Using such simple definitions of AS/AD shocks presents an immediate identification problem. We must identify four loadings (the responses of GDP and inflation to AS and AD shocks) but the covariance matrix of real activity and inflation shocks offers only three moments. In Section II, we lay out this identification problem and propose using higher-order moments to resolve it. We rely primarily on fourth-order moments, which are very significant in the data showing, economically and statistically, how identification is delivered by co-kurtosis moments. Along the way, we discuss the related econometrics literature (in particular, Lanne, Meitz and Saikkonen, 2017; and Lanne and Luoto, 2020).

Section III sets out the data we use, which include standard measures of real GDP growth and inflation, but also survey forecasts for these variables. It turns out that, in the context of a vector autoregression (VAR), survey data significantly improve forecasting power for real GDP growth and inflation. We also discuss our inference methodology, which uses a bootstrap procedure.

Section IV is the core of the paper. Our definitions of AS/AD shocks impose sign restrictions on the responses of inflation and real GDP to those shocks. There is a vast literature employing sign restrictions to aid identification in macroeconomic, mostly VAR, analyses (Faust, 1998; Uhlig, 1998, 2017; Canova and de Nicolo, 2002; Uhlig, 2017). Baumeister and Hamilton (2015) use sign restrictions to help achieve identification in a model of labor supply and demand in a Bayesian setting. Uhlig (2005) shows how to conduct inference in a VAR with sign restrictions imposed on the effects (measured using VAR impulse responses) of a monetary policy shock. Our setup here is most reminiscent of Uhlig's (2017) discussion of a price/quantity supply/demand identification problem. Consistent with Uhlig's proposals, we identify the sensitivity of reduced-form shocks to the structural shocks while assuming the structural shocks are mutually independent. Adding sign restrictions to information from the covariance matrix substantially restricts the set of feasible parameters, but it does not allow for complete identification. We show how to achieve identification using a simple method of moments estimator that exploits information in kurtosis and co-kurtosis moments of real activity and inflation. We account for sampling error in all the second- and fourth-order moments used and find that these moments provide quite sharp identification of the loadings of reduced-form shocks onto AS/AD shocks.

Small-sample concerns immediately jump to mind when using fourth-order moments in relatively small samples of macro data. Using extensive Monte Carlo and bootstrapping exercises, we show that the non-Gaussianities that we exploit in the data are robust to small sample issues. In addition, we analyze the small sample properties of our estimators and provide alternative confidence sets based on Anderson-Rubin tests (see Magnusson and Mavroeidis, 2009). We also verify the robustness of our findings to alternative VAR specifications. These include a "larger" VAR that includes several other information variables instead of survey expectations, a VAR with time-varying parameters, and a regime-switching VAR (Hamilton, 1989).

In Section V, we provide two validation exercises for our identification methodology. First, we show that our AD shocks affect output only temporarily while our AS disturbances have long-lasting effects on output. These two effects are consistent with the popular "long-run" restrictions suggested in the classic paper by Blanchard and Quah (1989), even though we do not impose them ex-ante. Second, we put our identification scheme to work to characterize the nature of recessions in the US historical record going back to 1968. Using our demand and supply shocks, we characterize the 1981-1982 and 2001 recessions as demand driven, whereas the 1969-1970, 1973-1975, and 1980 recessions are mostly supply driven. The results on the early recessions mostly corroborate earlier work by e.g. Gali (1992).

We then apply our work to study two notable episodes in recent US macroeconomic history: the Great Recession and the Covid-19 crisis. There is a debate on the origins of the Great Recession. Some researchers suggest that the event is well described as a large negative aggregate demand shock (see, e.g., Mian and Sufi, 2014; Mian, Rao and Sufi, 2013; and Benguria and Taylor, 2020). In contrast, Ireland (2011) and Mulligan (2012) stress the importance of supply shocks. We find that negative demand shocks contributed more to the Great Recession than supply shocks, though both AD and AS shocks were large.

We finally proceed to quantify the AS/AD composition of the Covid-19 recession and recovery, contributing to a rapidly developing literature studying the effects of the Covid-19 episode on macroeconomic activity. We estimate that the real GDP growth shocks for 2020:Q1 and 2020:Q2 were -1.63% (not at an annual rate) and an astounding -8.81%, respectively, and both were due in roughly equal parts to negative demand and supply shocks. These findings are consistent with del Rio-Chanona et al (2020), who provide estimates of "first-order" supply and demand shocks using micro data for the Covid-19 episode and also find both large demand and supply components. As the economy recovered in 2020:Q3-2021:Q1, with shocks to real GDP totaling a positive 4.63% (not at an annual rate), positive demand shocks accounted for essentially all of the recovery, with supply shocks contributing a small negative amount.

II Modeling and estimating aggregate supply and demand shocks

In this section, we outline how we identify and estimate aggregate supply and demand shocks by combining sign restrictions and higher order moments of macroeconomic data.

A *Defining and identifying aggregate supply and aggregate demand shocks*

Consider a bivariate system in real GDP growth, g_t, and inflation, π_t:

$$\begin{aligned} g_t &= E_{t-1}[g_t] + u_t^g \\ \pi_t &= E_{t-1}[\pi_t] + u_t^\pi \end{aligned} \tag{1}$$

where E_{t-1} denotes the expectation operator conditional on information available at time $(t-1)$. The variables u_t^g and u_t^π are reduced-form shocks. As described in Section IV, we use a VAR to identify these reduced-form shocks. Next, we model the reduced-form shocks as linear combinations of two structural shocks, labeled supply and demand, and denoted u_t^s and u_t^d, respectively:

$$\begin{bmatrix} u_t^\pi \\ u_t^g \end{bmatrix} = M \begin{bmatrix} u_t^s \\ u_t^d \end{bmatrix} \text{ where } M = \begin{bmatrix} -\sigma_{\pi s} & \sigma_{\pi d} \\ \sigma_{gs} & \sigma_{gd} \end{bmatrix} \tag{2}$$

The σ parameters in M are the loadings of the reduced-form shocks onto the supply and demand shocks. Our sign restrictions are that all of the σ parameters are positive. The first fundamental economic shock, u_t^s, is an aggregate supply shock, defined so that positive supply shocks increase g_t but decrease π_t. More generally, supply shocks move GDP growth and inflation in opposite directions. The second fundamental shock, u_t^d, is an aggregate demand shock, which moves GDP growth and inflation in the same direction, with positive demand shocks driving up both g_t and π_t. Supply and demand shocks are assumed to be independent and, without loss of generality, to have unit variance. Our use of sign restrictions is different from the common methodology in macroeconomics, pioneered by Faust (1998), Canova and De Nicolo (2002) and Uhlig (2005), which impose sign restrictions on impulse responses to aid identification. Uhlig (2017) presents a framework which uses simple sign restrictions on demand and supply for a generic good in the same manner that we propose, also assuming that the shocks are independent (see Assumption A.7 in Uhlig, 2017.)

Our definitions of supply and demand constitute a minimal set of assumptions that are consistent with many Keynesian structural models (Blanchard, 1989; Gali, 1992). Section V verifies, ex-post, whether the long-run effects of aggregate demand shocks are reflected mostly in prices and wages, not in output, and aggregate supply shocks, which include shocks to productivity, are more likely to have long-run effects on output, as in Blanchard and Quah (1989). Of course, these supply and demand shocks definitions do not necessarily correspond to demand and supply shocks in, say, a New Keynesian framework (see e.g. Woodford, 2003) or identified VARs in the Sims tradition (Sims, 1980).[1]

Returning to the identification problem, note that the sample covariance matrix of the reduced-form shocks from the bivariate system in equation (1) only yields three unique moments, but we need to identify four σ coefficients in equation (2) to extract the supply and demand shocks. In particular, the unconditional covariance matrix for inflation and growth shocks is:

$$Cov\left([u_t^{\pi}, u_t^{g}]\right) = \begin{bmatrix} \sigma_{\pi s}^2 + \sigma_{\pi d}^2 & -\sigma_{\pi s}\sigma_{gs} + \sigma_{\pi d}\sigma_{gd} \\ -\sigma_{\pi s}\sigma_{gs} + \sigma_{\pi d}\sigma_{gd} & \sigma_{gs}^2 + \sigma_{gd}^2 \end{bmatrix} \tag{3}$$

There are infinitely many combinations of the loadings in the vector $[\sigma_{\pi s}, \sigma_{\pi d}, \sigma_{gs}, \sigma_{gd}]$ that are consistent with the sample covariance matrix of inflation and growth shocks, with

[1] Furlanetto, Ravazzolo, and Sarferaz (2019) also employ similar sign restrictions to help identify demand and supply shocks in a larger VAR with more variables. We apply our methodology to a larger VAR in Section IV.E. Furthermore, in some models the "supply" shocks might move real activity and inflation in the same direction: see, for instance, news shocks in Cochrane (1994).

only a subset matching the sign restrictions. Thus, only a feasible set of loadings (still an infinitude), and not a unique vector can be identified based on sign restrictions alone, a situation sometimes referred to as "set identification."

Researchers have taken a variety of approaches to conduct inference when faced with set identification. Perhaps the most common approach is to impose additional "zero restrictions" by assuming that certain structural shocks do not affect some reduced-forms shocks contemporaneously. Cholesky decompositions of the covariance matrix are an example of this approach. That approach is not appropriate for our problem for the simple reason that it is not plausible that supply and demand shocks do not affect both output and inflation contemporaneously. An alternative approach was pioneered by Uhlig (2005) who samples on an equal weighted basis from all possible combinations of loadings that obey the sign restrictions. Uhlig (2017) emphasizes that a number of additional assumptions can be used to narrow down the set of parameters that both obey the sign restrictions and fit the covariance matrix. Our key contribution to this literature is to use higher-order moments to provide sharper inference relative to what can be accomplished with set identification alone.[2]

B Sharpening inference using higher-order moments

We next show how higher-order moments can provide useful identifying information. This effort is part of a growing econometric literature on using higher-order moments to achieve identification. Lanne, Meitz, and Saikkonen (2017) provide a general VAR framework with non-Gaussian innovations, and show how higher order moments help obtain identification. Estimation proceeds by maximum likelihood, whereas we propose a moments-based estimation, imposing minimal assumptions. Lanne and Luoto (2021), in a contemporaneous paper, also propose a moments-based estimation, but their paper is ideologically different

[2] A large recent literature complements sign restrictions with additional information to obtain identification. Kilian and Murphy (2012) show that sign restrictions are not sufficient to identify AS and AD shocks in a model focusing on oil shocks; Antolin-Diaz and Rubio-Ramirez (2018) use "narrative sign restrictions, " and Canova and Paustian (2011) compare different model specifications based on the sign of the impulse response functions instead of likelihood tests which might be non-robust.

– they propose an identification of uncorrelated shocks and then assign them an economic meaning. We reverse the order, first defining the shocks economically, then showing how to estimate them. This approach is more in line with the original Uhlig (2005) article.[3]

We focus on fourth-order moments because those turn out to be the most strongly statistically significant in our sample. In alternative empirical settings, skewness statistics may prove helpful. Here are the analytic expressions for the univariate excess kurtosis moment (denoted $x\kappa t$) of the two reduced-form shocks under the assumptions of the structural model equations:

$$\begin{aligned} x\kappa t\left(u^{\pi}, u^{\pi}, u^{\pi}, u^{\pi}\right) &= \frac{\sigma_{\pi s}^{4}\ x\kappa_{s} + \sigma_{\pi d}^{4}\ x\kappa_{d}}{var^{2}\left(u^{\pi}\right)} \\ x\kappa t\left(u^{g}, u^{g}, u^{g}, u^{g}\right) &= \frac{\sigma_{gs}^{4}\ x\kappa_{s} + \sigma_{gd}^{4}\ x\kappa_{d}}{var^{2}\left(u^{g}\right)} \end{aligned} \tag{4}$$

where $x\kappa_s$ and $x\kappa_d$ represent the unconditional excess kurtosis of supply and demand shocks, respectively, and var denotes variance. This equation and the subsequent equations are derived in Online Appendix A. Note that the assumptions of unit variance of the structural shocks and their independence are imposed. Unsurprisingly, the excess kurtosis moments of the reduced-form shocks are increasing functions of the excess kurtosis of the structural shocks. Adding just these two moments cannot resolve the identification problem because we must also estimate two additional parameters ($x\kappa_s$ and $x\kappa_d$). However, three additional cross-excess kurtosis moments are also available. The first is a "symmetric" cross kurtosis moment involving second powers of both inflation and growth, which can be written:

$$\begin{aligned} x\kappa t\left(u^{\pi}, u^{\pi}, u^{g}, u^{g}\right) &= \frac{cov\left(u_{\pi}^{2}, u_{g}^{2}\right)}{var\left(u^{\pi}\right) var\left(u^{g}\right)} - 2\rho^{2}\left(u^{\pi}, u^{g}\right) \\ &= \frac{\sigma_{\pi s}^{2}\sigma_{gs}^{2}\ x\kappa_{s} + \sigma_{\pi d}^{2}\sigma_{gd}^{2}\ x\kappa_{d}}{var\left(u^{\pi}\right) var\left(u^{g}\right)} \end{aligned} \tag{5}$$

where $\rho\left(u^{\pi}, u^{g}\right)$ denotes the correlation coefficient of shocks to inflation and growth. (Recall that $var\left(u^{\pi}\right)$, $var\left(u^{g}\right)$ and $\rho\left(u^{\pi}, u^{g}\right)$ are simple functions of the σ parameters in M). The top line shows that this moment measures the degree to which the squared values of

[3] There is an additional voluminous literature that exploits dynamic properties of shocks, including heteroskedasticity (mostly implying unconditional non-Gaussianities), to achieve identification (see e.g. Brunnermeier et al, 2021; Gourieroux et al, 2020).

the two shocks covary (in excess of what would be expected just because they have non-zero correlation, which is captured by the second term). If inflation and growth tend to be volatile or quiescent at the same time, this moment is positive. As shown by the second line, under our linear structure, excess kurtosis of the structural shocks contributes positively to this moment. This moment may aid identification because it requires no additional parameters under the maintained assumption that supply and demand shocks are independent (in particular that supply and demand shocks do not themselves have excess cross kurtosis).

In addition, two other fourth-order moments are available that may further sharpen our inference. These are the two "asymmetric" cross kurtosis moments, for instance:

$$x\kappa t\left(u^{\pi},u^{\pi},u^{\pi},u^{g}\right)=\frac{cov\left(u_{\pi}^{3},u_{g}\right)}{var\left(u^{\pi}\right)^{3/2}var\left(u^{g}\right)^{1/2}}-3\rho\left(u^{\pi},u^{g}\right) \tag{6}$$

A positive value for this moment means that, conditional on a positive (negative) growth shock, inflation shocks tend to be positively (negatively) skewed. Once again, the second term controls for what would be expected just because the series are correlated. The analytic expressions for this moment and its sibling, are

$$\begin{aligned} x\kappa t\left(u^{\pi},u^{\pi},u^{\pi},u^{g}\right) &= \frac{-\sigma_{\pi s}^{3}\sigma_{gs}\ x\kappa_{s}+\sigma_{\pi d}^{3}\sigma_{gd}\ x\kappa_{d}}{var^{3/2}\left(u^{\pi}\right)var^{1/2}\left(u^{g}\right)} \\ x\kappa t\left(u^{\pi},u^{g},u^{g},u^{g}\right) &= \frac{-\sigma_{\pi s}\sigma_{gs}^{3}\ x\kappa_{s}+\sigma_{\pi d}\sigma_{gd}^{3}\ x\kappa_{d}}{var^{1/2}\left(u^{\pi}\right)var^{3/2}\left(u^{g}\right)} \end{aligned} \tag{7}$$

Excess kurtosis of the supply shock has a negative effect on both of these moments. The intuition is clear: it is supply shocks that cause inflation and real growth to move in opposite directions. If the supply shock tends to take on extreme values, there is a propensity for inflation shocks to be positive (negative) when growth shocks are extremely negative (positive). An example is a large stagflation shock. However, if the excess kurtosis of demand shocks dominates, then inflation shocks tend to be positive (negative) when growth shocks are extremely positive (negative). As an example of this kind of dependence, the deeply negative readings on growth during the 2008-2009 Great Recession and the Covid-19 crisis were accompanied by negative inflation shocks. To see how these moments may help identify the σ parameters, suppose that the kurtosis of supply and demand shocks is similar. Then a high inflation sensitivity to supply shocks ($\sigma_{\pi s}$) relative to its sensitivity to demand shocks ($\sigma_{\pi d}$) lowers the co-kurtosis moment with inflation to the third power much more than it lowers the moment with GDP growth to the third power, and vice versa, all else equal. Note that fitting these asymmetric excess cross-kurtosis moments does not require estimating any additional parameters, given the independence of the structural shocks.

With the aid of these fourth-order moments we can, in principle, identify the loading parameters. This is true in a counting sense: we must estimate six parameters, $[\sigma_{\pi s}, \sigma_{\pi d}, \sigma_{gs}, \sigma_{gd}, x\kappa_s, x\kappa_d]$ and we have eight available moments. That said, the degree to which these moments lead to sharp inference of the σ parameters obviously depends on how precisely we can estimate them.

III Data and Inference

A Data

Our primary measures of reduced-form macroeconomic shocks are innovations to actual, real time (not revised) real GDP growth and inflation, as measured by the GDP deflator. These are gathered from the Real Time Dataset for Economists at the website of the Federal Reserve Bank of Philadelphia. We use real-time data releases for GDP and inflation, following Tulip (2005) and Faust and Wright (2013), who argue that it makes little sense to evaluate forecasts against final-revised data, which may not be available until many quarters later and because subsequent revisions often reflect definitional changes as opposed to additional data becoming available.[4] First-release data also have the advantage of not being susceptible to look-ahead bias in that future data cannot affect current measurement. We provide a robustness check using the final revised data in Section IV.E. We also use survey-based expectations of real GDP and GDP inflation from the Survey of Professional Forecasters as

[4]Ghysels et at (2018) note that forecasting bond returns is substantially more challenging when using real-time macroeconomic data as opposed to final revised data.

information variables. In particular, we use one-quarter ahead forecasts that are made 5 months before the calendar-end date of each target quarter. The use of first-release data is advantageous when used in conjunction with survey-based data because the surveys attempt to forecast GDP and inflation under the definitions and measurement procedures that were in place at the time.

Our data are quarterly, the units of both real GDP growth and inflation are log differences multiplied by 400, and the sample period is 1968Q4-2020Q1. We intentionally exclude data from 2020Q2 onwards because the astronomical values of the shocks to macroeconomic data during the height of the Covid-19 crisis tend to swamp other variation in the sample. We instead investigate the Covid-19 episode as an out-of-sample exercise.

B Inference

To identify shocks to real GDP growth and inflation, we use a VAR. Cognizant of small-sample biases, we rely on robust bootstrap procedures to conduct inference regarding our VAR parameters and other "downstream" parameters such as the second-and higher-order unconditional moments of the VAR residuals. We begin with "XY" bootstrapping. In particular, for a vector of jointly endogenous variables, Y_t, assuming a first-order VAR, we first sample from the joint data vector $[Y_t, Y_{t-1}]$ in blocks of 20 quarters. We creating 10,000 synthetic samples $[Y_t, Y_{t-1}]^{samp}$, each having length equal to that of our data. We estimate VAR parameters for each synthetic sample. For the parameters of the first-order VAR, these are just the coefficients from a regression of Y_t^{samp} onto Y_{t-1}^{samp}. In addition, we use each bootstrapped data set to calculate the downstream estimated parameters, such as the higher-order moments of the VAR residuals. We use the covariance matrix of parameter estimates across synthetic samples as the sampling covariance matrix of the parameters.

The method of moments procedure delivering $[\sigma_{\pi s}, \sigma_{\pi d}, \sigma_{gs}, \sigma_{gd}, x\kappa_s, x\kappa_d]$ requires additional bootstrap procedures to account for sampling error in the weighting matrix. These are described in Outline Appendix B. Section IV.D considers further finite sample analyzes of key test statistics and parameter distributions. We also establish that non-Gaussianities in the data are strong, rejecting the null of Gaussianity with a powerful test, taking into account finite sample issues.

IV Identification Results

A Estimating reduced-form shocks

Our VAR for real GDP growth and inflation can potentially include measures of survey-based expectations as information variables. To select the optimal specification in terms of lag length and to determine whether the data support inclusion of the information variables, we use the Bayesian AICc and BIC criteria.[5] The top panel of Table 1 presents results from these model-selection tests. We find that among the four specifications we test, the VAR with lag length 1 and including the information variables is optimal according to both criteria. In addition, for each specification we report standard Ljung-Box test p-values for first-order serial correlation in the residuals for inflation and GDP. The information variables help remove residual autocorrelation for the inflation residual. Concretely, our chosen specification is

$$Y_t = \mu + AY_{t-1} + u_t \tag{8}$$

where $Y_t = [\pi_t, g_t, S_t(\pi_{t+1}), S_t(g_{t+1})]$. $S_t(\pi_{t+1})$ and $S_t(g_{t+1})$ are survey expectations for inflation and real GDP growth from the Survey of Professional Forecasters. Survey forecasts have been shown to be very informative about future inflation (see Ang, Bekaert and Wei, 2007; and Faust and Wright, 2013). We view GDP and inflation forecasts as variables that summarize and rapidly react to reams of information in the economy, rendering the inflation and growth residuals orthogonal to past information in a parsimonious fashion. Our simple approach to append the expectations variables to the VAR also differs from a recent literature combining standard VARs with information in survey forecasts, usually employing Bayesian

[5] AICc refers to the AIC criterion corrected for small sample sizes, see Cavanaugh (1997).

methods (see Wright, 2012; Kruger, Clark and Ravazzolo, 2017; Tallman and Zaman, 2020; and Doh and Smith, 2022). If the surveys represent rational expectations, the coefficients on these variables should be one, but we let the data decide how much weight to put on the forecast variables.

The estimated coefficients from the feedback matrix, A, in the VAR are reported in Panel B of Table 1. Conditional expectations under the VAR for inflation load with coefficients of roughly 0.40 and 0.60 on past data and survey expectations, respectively. The VAR-based conditional expectation for GDP growth loads with coefficients of about 0.15 and 0.82 on past data and survey expectations, respectively. Clearly, survey-based expectations are very informative for forecasting inflation and real activity. Survey-based expectations are substantially persistent, with high autoregressive coefficients. In addition, higher realized inflation and higher realized growth both forecast increases in survey-based expectations for inflation. Survey-based expectations for GDP growth react negatively to higher realized inflation. All of these coefficients are statistically significant according to our bootstrapped standard errors.

Figure 1 plots the residuals from the chosen VAR for inflation and real GDP growth with NBER-dated recessions indicated with shading. As can be seen in the top panel, inflation shocks were positive during the recessions in the 1970s, but tended to be negative during recessions from the 1980s forward. Turning to the bottom panel, as expected, shocks to real GDP are negative during recessions throughout the sample. These patterns significantly inform our characterization of recessions as being either “supply-driven” or “demand-driven” in subsequent sections.

B Identification of structural shocks: step-by-step illustration

To identify the structural shocks for the VAR, we model u_t as:

$$u_t = \begin{bmatrix} M \\ M^i \end{bmatrix} \begin{bmatrix} u_t^s \\ u_t^d \end{bmatrix} + \Omega \varepsilon_t \qquad (9)$$

where, as before, u_t is a 4x1 vector of shocks from the VAR in equation (8); M has already been introduced in equation (2) as the 2x2 matrix that contains the loadings of inflation and real GDP growth shocks onto supply and demand shocks, and represents the key set of parameters that we wish to identify. M^i is a 2x2 matrix that contains the loadings of the two information variables in the VAR onto supply and demand shocks, and ε_t is a 2x1 vector of i.i.d. shocks that affect only the information variables in the VAR, and not inflation or real GDP growth. Therefore, Ω is a 4x2 matrix with zeros in the upper 2x2 block. To illustrate how higher-order moments sharpen identification of M, we first examine estimation of the system using second-order moments and sign restrictions alone.

1 *Sign restrictions and second-order moments*

Table 2, Panel A reports estimates of the three unconditional second-order moments that are available from the time series estimates of u_t^π and u_t^g. The unconditional standard deviations of u_t^g and u_t^π are fairly precisely estimated, but the correlation between the two is close to zero at -0.16 and not statistically significant using our bootstrapped standard errors. As discussed above, there are infinitely many possible combinations of the loadings in $[\sigma_{\pi s}, \sigma_{\pi d}, \sigma_{gs}, \sigma_{gd}]$ that are all consistent with the sample covariance matrix of inflation and growth shocks, and this under-identification occurs because only three second-order moments are available while we have four parameters to estimate.

The top panel of Figure 2 presents an illustration of all loadings (each loading is a 4-tuple) that match the sample second moments exactly, and satisfy the sign restrictions. To index the sets of loadings and plot them in a systematic manner, the units of the horizontal axis are $\sigma_{\pi s}$, plotted on the (positive) interval $[0, 0.3]$. As a result, the 45 degree line of asterisks shows values for $\sigma_{\pi s}$. To calculate the combinations that fit the second-order moments, we condition on each value for $\sigma_{\pi s}$. Next, using equation (3), we have the relation $\sigma_{\pi d}^2 = \pm \left(\widehat{var}\left(u_\pi\right) - \sigma_{\pi s}^2 \right)$ with the plus symbols plotting the positive solutions to this equation. Continuing, conditional on $\sigma_{\pi s}$ and $\sigma_{\pi d}$, the expressions in equation (3) for the covariance between inflation and growth and for the variance of growth yield a system with two quadratic equations for simultaneously determining σ_{gs} and σ_{gd}. That is, $[\sigma_{gs}, \sigma_{gd}]$ solves $\widehat{var}\,(u_g) = \sigma_{gs}^2 + \sigma_{gd}^2$ and $\widehat{cov}\,(u_g, u_\pi) = -\sigma_{\pi s}\sigma_{gs} + \sigma_{\pi d}\sigma_{gd}$, conditional on each pair $[\sigma_{\pi s}, \sigma_{\pi d}]$ and the sample statistics. These solutions for σ_{gs} and σ_{gd} are plotted as the black diamonds and triangles, respectively. The four lines thus represent the combinations of loadings that both fit the sign restrictions and match all the three second-order moments exactly. This illustrates the limited "set" or "interval" identification that sign restrictions provide.

2 *Adding information in higher-order moments*

Even though we only use fourth-order moments to aid in identification, for completeness, we report in Table 3, Panels B and C, estimates for both third- and fourth-order moments of the VAR residuals for inflation and real growth. Using our bootstrapped standard errors, there are no individually significant third-order moments. There are, however, three fourth-order moments that are individually significant. In particular, both u_t^π and u_t^g have significant excess univariate kurtosis, and the symmetric cross-kurtosis moment, denoted $x\kappa t\,(u^\pi, u^\pi, u^g, u^g)$ in the table, is positive and significant as well. This suggests that the volatilities of inflation and growth tend to increase in tandem. Both asymmetric cross kurtosis moments are negative, suggesting, for instance, that growth shocks tend to be more negatively skewed when inflation shocks are positive. Panel D of Table 3 reports tests due to Jarque and Bera (1980, "JB") for detection of non-Gaussianity. Section IV.D. shows these tests to have excellent small-sample properties in our setting. Both univariate JB tests for GDP growth and inflation as well as the bivariate JB test strongly reject the null of Gaussianity using bootstrapped p-values.

The bottom panel of Figure 2 illustrate how these higher-order moments can be used to narrow down the set identification provided in the top panel. The panel reports results from the following exercise. Conditional on each configuration of $[\sigma_{\pi s}, \sigma_{\pi d}, \sigma_{gs}, \sigma_{gd}]$ in the top panel (that is, each vertical slice of parameters in the top panel) all of which fit the second moments perfectly and obey the sign restrictions, we estimate $x\kappa_s$ and $x\kappa_d$ to fit the univariate kurtosis moments and the three cross kurtosis moment of inflation and real growth as well as possible. Then we judge which sets of loadings are most supported by the fourth-order moments using a goodness-of-fit statistic. Specifically, we use a classical minimum distance (CMD, see Wooldridge 2002 for a textbook treatment) setup in which we seek to match five statistics (the $x\kappa t$ moments in Table 3, Panel C) using two parameters ($x\kappa_s$ and $x\kappa_d$ introduced in equation (4)). Formally,

$$J\left(\sigma_{\pi s}, \sigma_{\pi d}, \sigma_{gs}, \sigma_{gd}\right) = \min_{[x\kappa_s, x\kappa_d]} \left\{ \left(\theta^{fitted} - \widehat{\theta}\right)' \widehat{W} \left(\theta^{fitted} - \widehat{\theta}\right) \right\} \tag{10}$$

where θ^{fitted} is a 5x1 vector containing fitted values for the fourth order moments while $\widehat{\theta}$ contains the sample counterparts from Table 3. As shown by equations (4),(5) and (7), θ^{fitted} is a function of each vector $[\widehat{\sigma}_{\pi s}, \widehat{\sigma}_{\pi d}, \widehat{\sigma}_{gs}, \widehat{\sigma}_{gd}]$ from the top panel of Figure 2 as well as the parameters to be estimated by CMD, $x\kappa_s$ and $x\kappa_d$. Because the system is overidentified, the estimation requires a weighting matrix, $\widehat{W}$. We use the inverse of the (bootstrapped) covariance matrix of $\widehat{\theta}$ so that the J-statistic from each CMD estimation is a goodness-of-fit test statistic which is asymptotically distributed as $\chi^2(3)$ distribution.[6] The horizontal line in the panel depicts the 5 percent critical value for a $\chi^2(3)$ distribution. The series of J statistics exhibit a single well that visibly narrows the set of plausible configurations of the loading vector. The vertical line at 0.23 highlights the minimum of the J-statistic, which corresponds to the values for the σ parameters at the corresponding vertical line in the upper panel. All values below the horizontal line fail to reject a perfect fit of the higher-order moments.

Of course, this is just an informal illustration of how higher-order moments can sharpen inference. While we could identify the parameters associated with the minimum $J-$statistic, this approach would insist that the second-order moments be fit perfectly in the first step, ignoring sampling error in those moments. This does not use all of the available information

[6] This ignores the sampling error in $[\widehat{\sigma}_{\pi s}, \widehat{\sigma}_{\pi d}, \widehat{\sigma}_{gs}, \widehat{\sigma}_{gd}]$. Section IV.D. examines the finite-sample distribution of the J-statistic for the full estimation.

efficiently, which may be particularly important in our case because one of the second-order moments, the unconditional correlation between inflation and real GDP shocks, is quite weak in the data, as discussed above.

C One step identification

A more efficient way to estimate the elements of M, $[\sigma_{\pi s}, \sigma_{\pi d}, \sigma_{gs}, \sigma_{gd}]$, is to use a one-step CMD estimation in which the three available second-order moments and the five available fourth-order moments are used to jointly to select the values for M that also obey the sign restrictions. Specifically, we fit the 8 statistics

$$\theta = \begin{bmatrix} \sigma\left(u^{\pi}\right), \sigma\left(u^{g}\right), \rho\left(u^{\pi}, u^{g}\right), x\kappa t\left(u^{\pi}, u^{\pi}, u^{\pi}, u^{\pi}\right), x\kappa t\left(u^{g}, u^{g}, u^{g}, u^{g}\right), \\ x\kappa t\left(u^{\pi}, u^{\pi}, u^{g}, u^{g}\right), x\kappa t\left(u^{\pi}, u^{\pi}, u^{\pi}, u^{g}\right), x\kappa t\left(u^{\pi}, u^{g}, u^{g}, u^{g}\right) \end{bmatrix}$$

reported in Table 3 using the 6 parameters, $[\sigma_{\pi s}, \sigma_{\pi d}, \sigma_{gs}, \sigma_{gd}, x\kappa_s, x\kappa_d]$. This is a CMD system with two degrees of overidentification. Formally, the CMD estimation solves the problem:

$$[\sigma_{\pi s}, \sigma_{\pi d}, \sigma_{gs}, \sigma_{gd}, \kappa_s, \kappa_d] = \arg\min \left\{ \left(\theta^{fitted} - \widehat{\theta}\right)' \widehat{W} \left(\theta^{fitted} - \widehat{\theta}\right) \right\}. \tag{11}$$

Once again, we use the inverse of the (bootstrapped) covariance matrix of the 8 statistics as the weighting matrix, $\widehat{W}$. Clearly, $\widehat{W}$ is an important input to the calculation of the 6 CMD parameters. Even though the covariance matrix on which $\widehat{W}$ is based, is calculated using a bootstrap procedure that should be relatively robust to small sample problems, the weighting matrix is still a source of sampling error of which we need to take proper account. To do so, we use a compound bootstrap, which is described in Online Appendix B.

The results of the CMD estimation with the compound bootstrap used to generate standard errors are reported in Table 3, Panel A. The estimates of the loading matrix, M, are listed in the top row. All loadings are highly statistically significant. Inflation shocks load (negatively, in accordance with the imposed sign restriction) onto supply shocks with an implied coefficient of -0.2432 and positively onto demand shocks with a coefficient of 0.2292. These relatively balanced loadings (given that supply and demand shocks both have unit variance by assumption), imply that variation in inflation shocks owes, on an unconditional basis, about equally to variation in supply and demand shocks. For real growth, the loading on supply shocks is 0.4974 and the loading on demand shocks is 0.3626. Thus, about 65 percent ($0.4974^2/\left(0.4974^2+0.3626^2\right)$) of variation in growth shocks arises from supply shocks on an unconditional basis. For comparison, the large symbols in the top panel of Figure 2 plot the fitted loadings from this one-step procedure against the set of loadings that obey the sign restrictions and fit the second moments exactly. Not surprisingly, the loadings from the one-step analysis do not lie exactly on those lines because the one-step analysis need not fit the second moments perfectly. In addition, they do not lie on the vertical line because they need not minimize the J-stat associated with the fourth-order moments alone.

As shown in the second block of Table 3, Panel A, both supply and demand shocks are estimated to have significant, positive excess kurtosis of around 2.5. For comparison, this value is roughly equivalent to the excess kurtosis of the distribution of monthly returns for the S&P 500 index, widely recignized for its "fat tails,"which is 2.8 for the same sample period.

Panels B and C of Table 3 report how well the CMD system is able to match the fitted statistics. For ease of comparison, the sample statistics that are used in the estimation and their standard errors are reproduced from Table 2. Reported in square brackets underneath are the fitted values from the CMD estimation. All statistics are matched within one standard error. The overidentification statistic, reported at the bottom of the table, is 1.10, so that the model is not rejected (p-value 0.58).

Figure 3 depicts the supply and demand shocks that are recovered from the one-step exercise. These estimates are easily obtained by inverting equation (2) using the parameter estimates for M from the one-step CMD estimation. Supply shocks are plotted in the top panel, demand shocks are in the bottom panel. Variation in supply shocks is clearly most prevalent in the early part of the sample. This is consistent with episodes of stagflation (e.g. as prompted by oil price shocks) that occurred in the 1970s and are generally associated with large supply disruptions. However, large negative supply shocks are also apparent during the Great Recession of 2008-2009. Demand shocks are prominently negative for all recessions since the 1980s, in particular during the "monetary policy experiment" from 1979-1980. They also show a sharp negative spike during the Great Recession. Still, there is an apparent secular decline in demand shock volatility over the sample period, potentially consistent with the Great Moderation.[7]

D Econometric Issues

Our identification procedure relies on a method of moments estimation, which features higher order moments, raising concerns about poor small sample properties of our estimators, and the possibility of weak identification. This section summarizes our extensive econometric analysis assuaging these concerns. To conserve space, we relegate technical details and detailed empirical results to the Online Appendix.

First, we analyze the small sample properties of various tests for non-Gaussianity of the inflation and GDP growth shocks through Monte Carlo analysis. We consider three different tests. The first test is a simple moment test of the 5 excess (co) kurtosis moments for inflation and GDP growth being zero; the second test is of the LM type, where we set up a GMM system to estimate those kurtosis moments under the null of normality and consider the test of the over-identifying restrictions. Finally, we consider the standard JB test and a multivariate extension, inspired by Zhou and Shao (2014).

Our findings rely on three Monte Carlo experiments. The first draws the structural shocks (the demand and supply shocks, and the shocks for the two expectations variables) from a joint i.i.d. Gaussian distribution, representing the null. To consider power, we use

[7] Angeletos et al. (2020) recently define a "main business cycle" (MBC) shock from a large VAR on 10 macro variables, representing the combination of the reduced from shocks that maximizes the explanatory power for real activity variables. We replicated the Angeletos et al (2020) analysis for a sample extending to 2020 (See the Online Appendix for technical details), and we find that the MBC shock loads positively and significantly on both of our AS and AD shocks, although together they only span about one-half of the variation in the MBC. This suggests that the MBC shock may be comprised, in part, of two very different and economically meaningful shocks, one related to AS and another one related to AD.

two data generating processes with empirically relevant non-Gaussian distributions for the structural shocks. In one experiment, we use the t-distribution, in another a mixture of normals distribution. The data generating processes match the excess kurtosis of the shocks. We then reconstitute data using the estimated structural parameters (the σ's) and the VAR parameters, delivering an artificial sample to which we apply our two-step estimation process, including the moments estimation.

We find that the Wald moment test severely over-rejects in small samples and the LM test lacks power. However, the JB tests have both good size and power properties. Using these tests, we find very strong rejections of the null of Gaussianity for both inflation and GDP growth shocks in our data. The estimated structural shocks themselves are also reliably non-Gaussian.

Second, while our standard errors rely on small sample distributions generated by a compound bootstrap procedure (see Online Appendix B), such distributions may be ill behaved. To further assess the empirical distribution of the parameter estimates in Table 3, we compare the bootstrapped distributions to Gaussian distributions of equal mean and standard deviation, finding five of the six distributions for the parameters to be well approximated by the corresponding Gaussian distribution. The distribution for σ_s^π has a "shoulder" to the left of the mode, presumably reflecting that some draws which over- or under-sample particular historical periods can generate lower values for this parameter. Overall, these empirical distributions bolster our confidence that the standard errors we report (and which reflect all sources of sampling error in our estimation procedure) are meaningful.

Third, we further examine tests that are robust to potential weak identification, which, in our context, would arise from the orthogonality conditions not being very sensitive to parameter changes. These tests are known as Anderson-Rubin tests (AR henceforth), but we rely on the treatment in Magnusson and Mavroeidis (2009). In our setting, the AR test is equivalent to our J-statistic for overidentification, as it does not depend on the Jacobian of the orthogonality conditions with respect to the parameter values. The asymptotic distribution for the J-statistic is $\chi^2(2)$. Using the two empirically relevant Monte Carlo experiments described above for non-Gaussianity tests, we derive the small sample distributions of the J-statistic in the two cases. The J-test over-rejects in both cases, with the 5% finite sample size being 20% for the t-distribution and 35% for the mixture of normals distribution data generating process. Because our test value in the data is 1.10, we are quite far away from a rejection under either the asymptotic, or the two small sample distributions.

Given a reliable small sample distribution for the over-identification test, we construct "robust" AR confidence intervals for our main parameters. That is, we identify the set of parameter vectors that does not prompt a rejection of the overidentification test, at a particular confidence level, which we set at 90%. The procedure and results are fully described in Online Appendix F. The AR confidence intervals are generally quite comparable to, albeit a bit wider than, our bootstrap confidence intervals. The only slight concern perhaps, is that the AR-based interval comprises, albeit marginally, zero for $x\kappa_s$. Therefore, the AR evidence does cast some doubt on the strength of the evidence in favor of fat tailed supply shocks.

E Robustness

Our VAR, which incorporates survey expectations, and only uses two macro variables, is non-standard relative to a large empirical macroeconomics literature. In this section, we show that our results are robust when considering more standard "large" VARs, and when accommodating more complex dynamics (in particular, regime switches and time variation in the parameters). Again, we relegate technical details and detailed empirical results (including multiple plots) to the Online Appendix.

First, we consider a "large" VAR with inflation, real GDP growth, the earnings-price ratio for the S&P 500 (defined as trailing four-quarter earnings per share divided by the current price per share), the federal funds rate, the 10-year Treasury yield minus the federal funds rate, and real consumption growth.[8] Note that this VAR does not incorporate survey expectations. For this robustness exercise, we employ exactly the same methodology as we did for the parsimonious VAR to infer VAR residuals, and demand and supply shocks. We find that the residuals for the baseline VAR and the large VAR are correlated 0.90 for both inflation and GDP growth residuals. As the table in Online Appendix section C.3 shows, the supply and demand loadings for both specifications are numerically and statistically very close. In addition, the resulting supply and demand shocks are 0.89 correlated.[9]

Our baseline VAR does not include a monetary policy variable, which may potentially undermine our identification of pure AD and AS shocks from inflation and real GDP innovations. However, the larger VAR includes the federal funds rate. In a VAR with quarterly data, the typical identification would order GDP growth and inflation before the monetary policy variable as they are more slowly moving. If the VAR captures the reduced form dynamics well, the monetary policy responses to shocks should be reflected in the feedback matrix (as monetary policy may have longer-lasting effects), and our AS/AD identification should not be contaminated by monetary policy effects. In our baseline VAR, the use of expectations variables may also be helpful. Monetary policy shocks likely affect our survey forecast variables within a quarter as forecasters should respond to monetary policy news quickly, and condition their forecasts on interest rates. These forecasts then help provide better predictions of future inflation and GDP growth and thus cleaner shocks going forward.

Second, given the Wold theorem, the linear VAR should be a good description of the data, provided there are no structural breaks or other forms of non-stationarity. Including the survey forecasts is helpful in this respect as they may partially capture slowly moving local means (see e.g. Kozicki and Taylor, 1998), which a linear system in observed variables cannot. Still, to verify the effect of potential structural changes, we estimate two alternative specifications to the base line VAR. In one alternative specification, we consider time-varying

[8] Canova and Ferroni (2021) explicitly discuss the econometric pitfalls when using a VAR that has fewer shocks than the true data generating process.

[9] We also verify that the long-term impulse responses, discussed in Section V, are quite similar.

parameters, modeling the VAR parameters as random walks. The correlation between the resulting residuals for GDP growth and inflation with those of the base line VAR are 0.97 and 0.94, respectively. In addition, we estimate a regime switching VAR version of the baseline VAR, where all parameters, including the conditional variance of the residuals, can switch regimes, with the regime variable a two state Markov chain with constant transition probabilities. Once again, the results are robust and the correlation between the resulting residuals for GDP growth and inflation with those of the baseline VAR are 0.98, respectively 0.99. These alternative models are estimated using Bayesian methods.

As a last robustness exercise, we re-ran our VAR with the most recently available "revised" data, rather than with the first -release data. The reduced-form VAR residuals for first-release inflation and real GDP growth have correlations with their final-revised counterparts of 0.62 and 0.73, respectively. Those imperfect correlations naturally feed through to the inverted supply and demand shocks (see the plot in Online Appendix G). The large (recessionary) shocks are largely the same for final-revised versus first-release data, but there are some notable deviations regarding the smaller quarter-to-quarter shocks. The correlation between first release and final-revised shocks for supply and demand are 0.55 and 0.79, respectively.

V Empirical Applications

We now use our identified aggregate supply and demand shocks in two empirical exercises. First, we ask whether the shocks satisfy the long-run Keynesian restrictions often imposed on structural VARs. Second, we formally characterize whether historical recessions are supply or demand driven, devoting special attention to the Great Recession and the Covid pandemic.

A Long-run effect of supply and demand shocks

In this section, we verify whether our "weak" but rather non-controversial short-term restrictions are consistent with the long-run Blanchard and Quah restrictions (1989). They associate supply shocks with shocks that have a permanent effect on output and demand shocks with shocks that have only a temporary effect on output, identified through "long-run" restrictions on impulse responses from simple VARs.[10] While such long-run restrictions are used frequently in empirical macroeconomics, they have been often criticized. Lippi and Reichlin (1993) show that long-run restrictions implicitly require making additional assumptions on the moving average representation of the VAR, which are not implied by theory. Faust and Leeper (1997) show that inference regarding such long-run restrictions can be quite unreliable, and propose to use short-run restrictions, relying on medium- to long-horizon restrictions only as an informal diagnostic. Baumeister and Hamilton (2015) propose to use the long-run restrictions as a prior belief in Bayesian setting to complement alternative restrictions. Given these criticisms, we focus on medium rather than long run restrictions, using a horizon of 20 quarters, as is common in more recent literature (see Francis et al, 2014.)

Returning to our baseline VAR model, we can use the estimates of the loading matrices, M and M^i, together with the feedback matrix, A, to calculate impulse response functions for the endogenous variables to one standard deviation supply and demand shocks.[11] In particular, the cumulative response H periods ahead, CIR_{t+H}, of variables in the VAR is

[10] They obtain full identification by also assuming that neither shock has a long-run effect on unemployment in a bivariate VAR with output and unemployment. Subsequent work adds price variables to the basic VAR (see, e.g. King, Plosser, Stock and Watson, 1991, and Bayoumi and Eichengreen, 1993.)

[11] We estimate $\widehat{M^i}$ by regressing the shocks to the infomation variables onto the estimated supply and demand shocks. This procedure results in loadings of shocks to $S_t\left(\pi_{t+1}\right)$ onto supply and demand shocks of -0.0213 and 0.0289, respectively. Similarly, shocks to $S_t\left(g_{t+1}\right)$ load onto supply and demand shocks with coefficients of 0.0911 and 0.0432, respectively. All of these loadings are statistically different from zero at the one percent level. Thus, survey expectations mimic the structural behavior of the realized macroeconomic data with respect to AS and AD shocks. The relative loading of $S_t\left(g_{t+1}\right)$ onto supply shocks is larger than for the actual GDP shocks, consistent with supply shocks having a prolonged effect on expected growth.

given by

$$CIR_{t+H} = \sum_{j=1}^{H} \widehat{A}^j \begin{bmatrix} \widehat{M} \\ \widehat{M}^i \end{bmatrix} \tag{12}$$

Results for the cumulative response (over 20 quarters) of (log) real GDP and the price level are reported in Table 4. The cumulative response of GDP to positive supply shocks is positive, at 1.21%, consistent with supply shocks persistently increasing the productive capacity of the economy, a feature of many Keynesian models. In contrast, demand shocks have no significant cumulative effect on the level of GDP, with an insignificant cumulative response of −0.05%, also consistent with the intuition and empirical results in Blanchard (1989). Demand shocks do have a persistent positive effect on the price level, estimated at 1.84% which is statistically significant, consistent with demand shocks causing a temporary increase in inflation, while supply shocks have a persistent negative effect on the price level, although that coefficient is not significant. Thus, our identification scheme for supply and demand shocks, while imposing very little economic structure, produces medium-term dynamics consistent with those in popular Keynesian models identified through a more complex set of restrictions.

B How supply and demand shocks drive U.S. recessions

A key application of our framework is to quantify the extent to which historical U.S. recessions were driven by supply versus demand shocks, the subject of a large literature. Table 5 presents results of this exercise. For each NBER-defined recession since 1970, we estimate the contribution of supply and demand shocks to the downturn in two different ways. Under the first methodology, we simply calculate the sum of supply (and, separately demand) shocks over the recession period and use the estimated loading matrix, $\widehat{M}$, to calculate the effect of the total supply and demand shocks on real GDP. These results are presented in the left columns of Table 5. For example, during the 1970:Q1-1970:Q4 recession, GDP decreases by 1.14% out of which 1.47% is attributed to a negative AS shock and 0.33% to a positive AD shock (-1.47%+0.33%=-1.14%). We also report 95% confidence intervals for the point estimates in this table, which reflect the sampling uncertainty of $\widehat{M}$.[12] The recessions beginning in 1970 and 1973 featured large, statistically significant, negative supply shocks with mild demand shock components that are actually positive. The 1980 recession appears to have witnessed only small and insignificant supply and demand components. The more severe 1981 recession is attributed more to demand than supply shocks, consistent with a monetary policy-induced recession, although both types of shocks were important, statistically significant, drivers. The 1990 recession features both significantly negative demand and supply shocks with supply shocks being more important. This pattern is reversed for the 2008-2009 Great Recession. The milder 2001 recession exhibited negative demand shocks, which are statistically significant, but the supply shocks were positive.

The second method that we employ calculates "aggregated shocks," recognizing that the effects of shocks may be larger or smaller than their simple sum implies, under the dynamics implied by the estimated VAR. In particular, we define aggregated GDP shocks by computing

$$u_{t+N}^{g,agg} = \sum_{i=1}^{i+N} g_{t+i} - E_{t-1} \sum_{i=1}^{i+N} g_{t+i} \tag{13}$$

which measure the cumulative shortfall of real GDP (over the N periods of each recession) relative to expectations under the VAR on the eve of the recessions. Aggregate shocks for inflation, $u_{t+N}^{\pi,agg}$, are defined analogously. The components of these aggregated shocks that are due to supply and demand are calculated by multiplying $u_{t+N}^{g,agg}$ and $u_{t+N}^{\pi,agg}$ (which are functions of the residuals u_{t+i}^{g} and u_{t+i}^{π} $i = 1, ..., N$) by $\widehat{M}^{-1}$. Finally, supply and demand shocks are rescaled by σ_{gs} and σ_{gd} to put the results into units of GDP growth. The results of this exercise are reported in the right two columns of Table 5. We also report 95% confidence intervals for the point estimates in this part of the table, now reflecting sampling uncertainty of $\widehat{M}$ as well as the VAR parameters. The estimated cumulative impact of the shocks is much

[12] To do so, we calculate the estimated supply and demand effects for each recession using estimates of M from each of 10,000 bootstrapped samples in the "outer bootstrap" described in Online Appendix B.

larger than implied by the first methodology, consistent with persistence implied by the A matrix. Still, a consistent picture emerges: recessions up through the one that began in 1980 were predominantly supply driven. Starting with the 1981 recession, most recessions were more demand driven, again with the 1990 recession as an exception. The Great Recession is estimated to have both a large supply and a large demand component, with the demand component dominating.

For the first four recessions, we can compare our results with results in the classic Gali (1992) article. His set-up is more complex than ours in that he distinguishes not only between supply and demand (his "IS") shocks, but also money supply and demand shocks. Gali provides variance decompositions for the first 4 recessions in terms of the different shocks. The seventies recessions are predominantly supply driven (with the 1973 recession 78% driven by supply shocks); and the eighties recessions are predominantly demand driven (although money supply shocks account for 38% of the 1981-1982 recession). Thus, our results are broadly consistent with Gali's.

C The Great Recession

The Great Recession of 2008-2009 deserves special attention because it generated a huge literature on its origins and economic effects. While the crisis was financial in nature, it is still of substantial economic interest to ascertain whether it was primarily supply or demand driven. This is useful for judging the appropriateness of policy responses (e.g. fiscal or monetary stimulus to counteract adverse demand shocks; tax cuts or structural reform to mitigate adverse supply shocks). A great many articles address the question using a large variety of economic models and empirical techniques. Mian and Sufi (2014) use micro data to argue that a deterioration in household balance sheets, and the accompanying lower consumption, ultimately was the main determinant of the sharp decline in aggregate U.S. employment between 2007 and 2009. Benguria and Taylor (2020) study historical financial crises including the Great Recession from the perspective of a small open economy model to identify whether crises are demand or supply driven. They conclude that financial crises, including the Great Recession, predominantly represent a negative shock to demand.

While Mulligan (2011, 2012) does not couch his analysis in terms of aggregate demand and supply shocks, his analysis of labor market distortions during the Great Recession points to supply factors, such as government regulations reducing incentives to work, and higher real wages contributing to lower labor hours. Supply shocks then explain why the Great Recession lasted so long. Ireland (2011) uses a New Keynesian micro-founded model to conclude that the Great Recession was caused by a mixture of supply and demand shocks. Ireland's model features 4 structural shocks and he identifies preference and monetary policy shocks as "demand shocks," and technology and cost push shocks as "supply shocks," exactly because they satisfy our definition in equation (2). Note that in many DSGE models, certain structural shocks may have both the supply and demand effects we describe depending on the parameter values.

Relative to this literature, our approach is minimalistic, as we only use GDP growth and inflation data (in addition to survey data), and weak sign restrictions on their responses to demand and supply shocks. Yet, this suffices to obtain identification, and delivers results consistent with the results from Ireland's (2011) model. While aggregate demand shocks dominate, cumulating to -3.01%, aggregate supply shocks are also important, cumulating to -2.39%. Using our first (second) methodology, demand shocks account for 59.1% (55.1%) of the total negative shock. Clearly, if the Great Recession was driven by a pure demand shock, it would have been more deflationary. In terms of magnitudes, the Great Recession shocks were of the same order of magnitude as the shocks for the 1981-1982 recession, with the supply (demand) shocks more (less) severe.

D The Covid Crisis

Our final application is to decompose the shocks to real GDP growth during the Covid-19 pandemic and its aftermath. To do so, we estimate (out-of-sample) residuals from the

VAR for 2020:Q1 through 2021:Q2. We then uncover the supply and demand shocks using equation (2), and scale them by σ_{gs} and σ_{gd}, respectively, to estimate the effect on GDP. The results are reported in Table 5. Once again, 95% confidence intervals are also reported to reflect sampling uncertainty in $\widehat{M}$. The shock to GDP in 2020:Q1 was very large by historical standards (-1.63%) and was driven in roughly equal proportions by supply and demand shocks. (All growth rates cited in this section are quarterly, not annualized.) The shock in 2020:Q2 was unprecedented in size, with GDP unexpectedly declining by -8.81%. According to our decomposition, demand and supply shocks contributed about equally to the decline. That both kinds of shocks affected the macroeconomy is potentially consistent with Guerrieri et al (2020) who argue that supply shocks may have fed back to spur additional demand shocks during the Covid-19 pandemic (although such interplay between supply and demand shocks would violate our assumption of independence).

Table 6 also shows the shocks during the initial phase of the economic recovery, from 2020:Q3 to 2021:Q1. There was a large, positive shock to real GDP in 2020:Q3 relative to expectations under the VAR. (As a reminder, expectations for real GDP growth under the VAR load strongly on survey forecasts, suggesting that SPF forecasters were also substantially surprised by the speed of the recovery.) The net positive shock of 4.95% is estimated to consist of a positive contribution of 6.22% from demand, and a contribution of -1.27% from supply. The negative supply shock during the recovery phase could reflect that supply bottlenecks continued to worsen even as the economy recovered, which is plausible in light of the uneven pace of reopening across sectors and geographic regions. In 2020:Q4, demand and supply shocks were small and insignificant, as was the total shock to real GDP. Finally, in 2021:Q1 there was a positive shock of 0.85% to real GDP growth, which reflected a positive shock to demand. Thus, through the end of this sample, the recovery was hampered by a lack of positive supply shocks.

The 2021-2022 macroeconomic landscape was dominated by an inflationary surge against the backdrop of both large adverse aggregate supply shocks (a war between Russia and the Ukraine; continued lockdowns in China) and large Covid-related and post-Covid fiscal stimulus packages. Monetary policy started to tighten and increase interest rates. In such an economic environment, the relative importance of AS and AD shocks is obviously an important ingredient in the debates among economists, policy makers and market observers about the right policy response.

VI Conclusion

In this article, we show how to combine minimal short-term sign restrictions with intuitive higher-order moments of inflation and real GDP growth to identify aggregate demand and aggregate supply shocks for the US economy. Methodologically, our set-up can be extended to other structural VARs, as long statistically significant higher order moments are available to aid identification.[13] In a robustness section, we show that our results survive in various alternative specifications, including a "large" VAR, a regime -switching VAR and a VAR with time-varying parameters.

Empirically, we show that the methodology generates AS and AD shocks that satisfy popular long-run restrictions used in the Keynesian literature to identify structural VARs. We also find that the relative importance of the AS and AD shocks in the seventies and eighties recessions are in line with earlier work by Gali (1992). We attribute the Great Recession slightly more to AD shocks than to AS shocks, but the latter are still quite important, accounting for more than 40% of the total negative shocks. A similar picture emerges for the Covid-19 crisis. While the crisis precipitated unprecedented large negative supply and demand shocks, the recovery starting in 2020:Q3 has been surprisingly strong with respect to demand shocks. In contrast, supply shocks have continued to be negative, on net, during the recovery. This is consistent with additional supply bottlenecks continuing to hamper economic activity, and putting upward pressure on inflation.

[13] See Bekaert, Engstrom, and Ermolov, 2021, for an application to the US term structure.

References

Ang, A., Bekaert, G., and M. Wei, 2007, "Do Macro Variables, Asset Markets, or Surveys Forecast Inflation Better?" Journal of Monetary Economics, 4, 2007, ppg. 1163-1212.

Angeletos, G., Collard, F. and H. Dellas, 2020, "Business-Cycle Anatomy," American Economic Review, vol. 110, no. 10, ppg. 3030-3070.

Antolin-Diaz, J. and J. Rubio-Ramirez, 2018, "Narrative Sign Restrictions for SVARs," American Economic Review, vol. 108, iss. 10, ppg. 2802-29.

Baumeister, C. and J. Hamilton, 2015, "Sign Restrictions, Structural Vector Autoregressions, and Useful Prior Information," Econometrica 83(5), pp. 1963–1999.

Bayoumi, T. and B. Eichengreen, 1993, "Shocking Aspects of European Monetary Integration," Adjustment and Growth in the European Union (F. Torres and F. Giavazzi eds.), pp. 193-235.

Bekaert, G., Engstrom, E., and A. Ermolov, 2021, "Macro Risks and the Term Structure of Interest Rates," Journal of Financial Economics, forthcoming.

Benguria, F. and A. Taylor, 2020, "After the Panic: Are Financial Crises Demand or Supply Shocks: Evidence from International Trade," American Economic Review: Insights, 2(4), pp. 509-526.

Blanchard, O., and D. Quah, 1989, "The Dynamic Effects of Aggregate Demand and Supply Disturbances," American Economic Review, 79, pp. 655-673.

Blanchard, O., 1989, "A Traditional Interpretation of Macroeconomic Fluctuations," American Economic Review, 79, pp. 1146-1164.

Brunnermeier, M., Palia, D., Sastry, K., and C. Sims, 2021, "Feedbacks: Financial Markets and Economic Activity," American Economic Review, vol. 111, no. 6, ppg. 1845-1879.

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

Canova, F. and F. Ferroni, 2021, "Mind the Gap! Stylized Dynamic Facts and Structural Models," American Economic Journal: Macroeconomics, vol. 14, iss. 4, ppg. 104-153.

Canova, F. and M. Paustian, 2011, "Business Cycle Measurement with some Theory," Journal of Monetary Economics, vol. 58, iss. 4, ppg 345-361.

Cavanaugh, Joseph E., "Unifying the derivations for the Akaike and corrected Akaike information criteria," Statistics and Probability Letters, 33(2), pp. 201-208.

Cochrane, J., 1994, "Shocks," Journal of Monetary Economics, 41, pp. 295-364.

del Rio-Chanona, R. , Mealy, M., Penny, Pichler, P., LaFond, A., Fracois and J. Farmer, 2020, "Supply and demand shocks in the COVID-19 pandemic: an industry and occupation perspective," Oxford Review of Economic Policy, 36, pp. S94-S137.

Doh, T. and L. Smith, 2022, "A New Approach to Integrating Expectations into VAR Models," Journal of Monetary Economics, forthcoming.

Faust, J., 1998, "The Robustness of Identified VAR Conclusions About Money," Carnegie-Rochester Conference Series on Public Policy 49, 207–244.

Faust, J. and E. Leeper, 1997, “When do Long-Run Identifying Restrictions Give Reliable Results?” Journal of Business and Economic Statistics, 15(3), pp. 345-353.

Faust, J. and J. Wright, 2013, “Forecasting Inflation,” Handbook of Economic Forecasting, 2, pp. 2-56.

Francis, N., Owyang, M., Rousch, J., and R. DiCecio, 2014, “A Flexible Finite-Horizon Alternative to Long-Run Restrictions with an Application to Technology Shocks,” Review of Economics and Statistics, vol. 96, iss. 4, ppg. 638-647.

Furlanetto, F., Ravazzolo, F., and S. Sarferaz, “Identification of Financial Factors in Economic Fluctiations,” The Economic Journal, vol. 129, iss. 617, ppg. 311-337.

Gali, J., 1992, “How Well Does the IS-LM Model Fit Postwar U.S. Data?” Quarterly Journal of Economics, 107, pp. 709-738.

Ghysels, E., Horan, C., and E. Moench, “Forecasting through the rearview mirror: Data Revisions and Bond Return Predictability,” Review of Financial Studies, 31(2), pp. 678-714.

Gourieroux C. J. Jasiak, 2022, "Semi-Parametric Estimation of Noncausal Vector Autoregression", Journal of Econometrics, forthcoming.

Guerrieri, V., Guido, L., Ludwig, S., and I. Werning, 2020, “Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages?” Working paper.

Hamilton, J., 1989, “A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle,” Econometrica, 57, pp. 357-384.

Ireland, P., 2011, “A New Keynesian Perspective on the Great Recession, Journal of Money, Credit and Banking,” 43, pp. 31-54.

Jarque, C. and A. Bera, 1980, “Efficient tests for Normality, Homoscedasticity and Serial Independence of Regression Residuals,” Economics Letters. 6 (3): 255–259.

Kilian, L. and D. Murphy, 2012, “Why Agnostic Sign Restrictions are not Enough: Understanding the Dynamics of Oil Market VAR Models,” Journal of the European Economic Association, Vol. 10, No. 5 (October 2012), pp. 1166-1188.

King, R., Plosser, C., Stock, J., and M. Watson, 1991, “Stochastic Trends and Economic Fluctuations,” American Economic Review, Sept. 1991, pp. 819-840.

Kozicki, S. and P. Tinsley, 1998, “Moving Endpoints and the Internal Consistency of Agents’ Ex Ante Forecasts,” Computational Economics, vol. 11, ppg. 21-40.

Kruger, F., Clark, T., and F. Ravazzolo, 2017, “Using Entropic Tilting to Combine BVAR Forecasts with External Nowcasts,” Journal of Business and Economic Statistics, vol. 35, no. 3, ppg. 470-485.

Lanne, M., and J. Luoto (2020). Identification of Economic Shocks by Inequality Constraints in Bayesian Structural Vector Autoregression, Oxford Bulletin of Economics and Statistics 82, 425 – 452.

Lanne, M., M., Meitz, and P. Saikkonen, 2017, “Identification and Estimation of Non-Gaussian Structural Vector Autoregressions,” Journal of Econometrics, 196, pp. 288-304.

Lippi, M. and L. Reichlin, 1993, “The Dynamic Effects of Aggregate Demand and Supply Disturbances: Comment,” American Economic Review, 83(3), pp. 644-652.

Magnusson, L., and S. Mavroeidis, “Identification-Robust Minimum Distance Estimation of the New Keynesian Phillips Curve,” 2010, Journal of Money, Credit and Banking, vol. 42, No. 2/3 (March-April 2010), pp. 465-481

Mian, A., Rao, K., and A. Sufi, 2013, “Household Balance Sheets, Consumption, and the Economic Slump,” Quarterly Journal of Economics, July, pp. 1687-1726.

Mian, A. and A. Sufi, 2014, “What Explains the 2007-2009 Drop in Employment?” Econometrica, 82, pp. 2197-2223.

Mulligan, C., 2011. “Rising Labor Productivity during the 2008-9 Recession,” NBER Working Papers 17584, National Bureau of Economic Research, Inc.

Mulligan, C., 2012, “The Redistribution Recession: How Labor Market Distortions Contracted the Economy,” Oxford University Press.

Shapiro, M. and M. Watson, 1988, “Sources of Business Cycle Fluctuations,” NBER Macroeconomics Annual 1988, Volume 3, pp. 111-156.

Sims, C., 1980, “Macroeconomics and Reality,” Econometrica, 48, pp. 1-48.

Tallman, E. and S. Zaman, 2020, “Combining Survey Long-Run Forecasts and Nowcasts with BVAR Forecasts Using Relative Entropy,” International Journal of Forecasting, vol 36, iss. 2, ppg. 373-398.

Tulip, P., 2005. “Has output become more predictable? Changes in Greenbook forecast accuracy," Finance and Economics Discussion Series 2005-31, Board of Governors of the Federal Reserve System (U.S.).

Uhlig, H., 1998, “The Robustness of Identified VAR Conclusions About Money: A Comment,” Carnegie-Rochester Conference Series on Public Policy 49, 245–263.

Uhlig, H., 2005, “What Are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure,” Journal of Monetary Economics, 52, pp. 381-419.

Uhlig, H., 2017, “Shocks, Sign Restrictions and Identification,” In Honore, B., Pakes, A., Piazzesi, M., & Samuelson, L. (eds.), Advances in Economics and Econometrics: Eleventh World Congress (Econometric Society Monographs), Cambridge University Press (2017), 95-127.

Woodford, M., 2003, “Interest and Prices: Foundations of a Theory of Monetary Policy,” Princeton University Press.

Wooldridge, Jeffrey, 2002, “Econometric Analysis of Cross Section and Panel Data,” MIT Press.

Zhou, M. and Y. Shao, 2014, “A Powerful Test for Multivariate Normality,” Journal of Applied Statistics, vol. 41, iss. 2, ppg. 351-363.

Table 1: VAR Specification Tests and Parameter Estimates

Panel A: Specification tests

		Info Criteria		Q pvals	
lags	info vars	AICc	BIC	u_t^π	u_t^g
1	no	560.2	590.1	0.000	0.852
2	no	544.4	587.6	0.168	0.966
1	yes	490.6	533.8	0.836	0.222
2	yes	496.7	553.2	0.563	0.563

Panel B: Parameter estimates

	π_{t-1}	g_{t-1}	$S_{t-1}(\pi_t)$	$S_{t-1}(g_t)$
π_t	0.4104	0.0287	0.6133	−0.0116
	(0.1312)	(0.0519)	(0.1225)	(0.0727)
g_t	−0.3420	0.1513	0.2868	0.8190
	(0.2202)	(0.0963)	(0.2104)	(0.2243)
$S_t(\pi_{t+1})$	0.2355	0.0579	0.7285	−0.0029
	(0.0440)	(0.0178)	(0.0419)	(0.0240)
$S_t(g_{t+1})$	−0.1975	−0.0558	0.1749	0.8230
	(0.0969)	(0.0557)	(0.1252)	(0.0993)

Notes: Data are quarterly 1968Q4-2020Q1. VAR variables include the log growth rate of real GDP, the log growth grate of the GDP deflator, and, for specifications including information variables, one-quarter ahead forecasts for real GDP growth and inflation of the GDP price deflator, both from the Survey of Professional forecasters. In Panel A, results under "Info Criteria" columns are Bayesian information criteria for various VAR specifications. AICc and BIC refer to the Akaike and Schwarz criteria, respectively. To calculate the AICc and BIC criteria, for all specifications we use the residuals and parameter count for just the first two elements of the VAR: GDP growth and inflation. Results under "Q" columns are p-values for Ljung Box Q tests of autocorrelation for the VAR's inflation residual (u_t^π) and its GDP growth residual (u_t^g) at the 1-quarter horizon. In Panel B, we report estimates for the matrix A in the estimated VAR equation $Y_{t+1} = \mu + AY_t + u_t$. OLS point estimates are reported with bootstrapped standard errors in parentheses.

Table 2: Inflation and GDP growth moments

Panel A: Second-order moments

$\sigma(u^\pi)$	$\sigma(u^g)$	$\rho(u^\pi, u^g)$
0.3181	0.5829	−0.1654
(0.0402)	(0.0855)	(0.0981)

Panel B: Scaled (co-)skewness

$skw(u^\pi, u^\pi, u^\pi)$	$skw(u^g, u^g, u^g)$	$skw(u^\pi, u^\pi, u^g)$	$skw(u^\pi, u^g, u^g)$
0.2968	0.1016	−0.3934	0.2325
(0.3528)	(0.2152)	(0.2499)	(0.2094)

Panel C: Scaled excess (co-)kurtosis

$xkt(u^\pi, u^\pi, u^\pi, u^\pi)$	$xkt(u^g, u^g, u^g, u^g)$	
1.4610	1.7987	
(0.5519)	(0.8036)	
$xkt(u^\pi, u^\pi, u^g, u^g)$	$xkt(u^\pi, u^\pi, u^\pi, u^g)$	$xkt(u^\pi, u^g, u^g, u^g)$
1.2212	−0.3382	−0.6680
(0.4804)	(0.7088)	(0.4656)

Panel D: Jarque-Bera Tests for non-Gaussianity (p-vals)

$JB(u_t^\pi)$	< 0.001
$JB(u_t^g)$	< 0.001
$JB(u_t^\pi, u_t^g)$	< 0.001

Notes: Data are quarterly 1968Q4-2020Q1. VAR setup is $Y_{t+1} = \mu + AY_t + u_t$. Variables include the log growth rate of real GDP, the log growth grate of the GDP deflator, and one-quarter ahead forecasts for real GDP growth and inflation of the GDP price deflator, both from the Survey of Professional Forecasters. The time series u_t^π and u_t^g are the residuals for inflation and real GDP growth. The notation $\sigma(u^\pi)$ denotes the unconditional standard deviation of u_t^π, and $\rho(u^\pi, u^g)$ denotes the unconditional correlation between the two time series. The notation for third and fourth-order moments is as follows: $skw(u^\pi, u^\pi, u^\pi)$ denotes the unconditional, scaled third-order moment $E\left[(u^\pi)^3\right] / E\left[(u^\pi)^2\right]^{3/2}$ the symbol xkt notation is used for excess kurtosis moments. All indicated unconditional, scaled, third- and fourth-order moments are in excess of what would be implied by the Gaussian distribution. Bootstrapped standard errors are reported in parentheses. Panel D reports the p-values for univariate and a bivariate Jarques-Bera test for Gaussianity.

Table 3: One-step CMD Identification

Panel A: Parameter estimates

$\widehat{\sigma}_{\pi s}$	$\widehat{\sigma}_{\pi d}$	$\widehat{\sigma}_{gs}$	$\widehat{\sigma}_{gd}$
0.2423	0.2292	0.4974	0.3626
(0.0634)	(0.0333)	(0.0792)	(0.0853)
$\widehat{\kappa_s}$	$\widehat{\kappa_d}$		
2.4537	2.4328		
(0.5924)	(0.6718)		

Panel B: Fitted second-order moments

$\sigma(u^\pi)$	$\sigma(u^g)$	$\rho(u^\pi, u^g)$
0.3181	0.5829	-0.1654
(0.0402)	(0.0855)	(0.0981)
[0.3344]	[0.6177]	[-0.1840]

Panel C: Fitted fourth-order moments

$xkt(u^\pi, u^\pi, u^\pi, u^\pi)$	$xkt(u^g, u^g, u^g, u^g)$	
1.4610	1.7987	
(0.5519)	(0.8036)	
[1.2260]	[1.3392]	
$xkt(u^\pi, u^\pi, u^g, u^g)$	$xkt(u^\pi, u^\pi, u^\pi, u^g)$	$xkt(u^\pi, u^g, u^g, u^g)$
1.2212	-0.3382	-0.6680
(0.4804)	(0.7088)	(0.4656)
[1.2443]	[-0.2954]	[-0.5990]

$Jstat$	$p-value$
1.10	(0.58)

Notes: Data are quarterly 1968Q4-2020Q1. In Panel A, the parameter estimates from the one-step CMD exercise described in the text are reported. In Panel B, the times series u_t^π and u_t^g are the residuals for inflation and real GDP growth from the optimal VAR identified in the previous section. The notation $\sigma(u^\pi)$ denotes the unconditional standard deviation of u_t^π, and $\rho(u^\pi, u^g)$ denotes the unconditional correlation between the two time series. In Panel C, the notation for fourth-order moments is as follows: $xkt(u^\pi, u^\pi, u^g, u^g)$ denotes the unconditional excess cross-kurtosis moment $\frac{cov\left(u_\pi^2, u_g^2\right)}{var(u^\pi)var(u^g)} - \rho^2(u^\pi, u^g)$ and other moments similarly denote corresponding unconditional, scaled, fourth-order moments in excess of what would be implied by the Gaussian distribution. Bootstrapped standard errors are reported in parentheses. Numbers in square brackets denote the fitted values for each statistic under the one-step estimation procedure described in the text. The J-statistic and p-value are for the overidentification test of the CMD system.

Table 4: Cumulative impulse response functions

shock	real GDP level	price level
demand	-0.05%	$1.84\%^{***}$
	(54.33%)	(0.00%)
supply	$1.21\%^{***}$	-0.89%
	(0.01%)	(94.01%)

Notes: Cumulative VAR Impulse Responses of Real GDP and Aggregate Price Level to One Standard Deviation Demand and Supply Shocks. The sample is quarterly 1968Q4-2019Q4. The VAR setup is $Y_{t+1} = \mu + AY_t + \begin{bmatrix} M \\ M^i \end{bmatrix} \begin{bmatrix} u_t^s \\ u_t^d \end{bmatrix} + \varepsilon_t$. where Y_t is the vector of GDP growth, inflation and their one period ahead Survey of Professional Forecasters expectations, and u_t^s and u_t^d are pre-estimated supply and demand shocks, respectively. Impulse responses are cumulative over 20 quarters including the quarter where the shock happened. Numbers in parentheses are probabilities that the impulse response is less than 0 obtained from 10,000 block-bootstrap samples of historical length. The asterisks, ***, correspond to statistical significance at the 1% level.

Table 5: Decomposition of recession shocks

	individual shocks		aggregate shock	
NBER recession	demand	supply	demand	supply
1970:Q1-1970:Q4	0.46	−0.86	0.33	−1.47
	(0.09, 0.85)	(−1.48, −0.46)	(−0.10, 0.80)	(−2.39, −0.47)
1973:Q4-1975:Q1	0.19	−1.08	3.25	−8.96
	(−1.53, 0.93)	(−5.11, −0.60)	(1.76, 4.74)	(−11.25, −6.33)
1980:Q1-1980:Q2	0.11	−0.29	0.45	−1.35
	(−0.27, 0.73)	(−1.08, 1.15)	(0.13, 0.96)	(−2.34, 0.28)
1981:Q3-1982:Q4	−2.40	−1.08	−4.11	−1.19
	(−3.23, −0.57)	(−2.89, 0.09)	(−5.87, −0.84)	(−4.25, 1.77)
1990:Q4-1991:Q1	−0.48	−0.85	−0.72	−1.02
	(−0.74, −0.05)	(−1.39, −0.25)	(−1.05, −0.11)	(−1.75, −0.49)
2001:Q2-2001:Q4	−0.35	0.64	−1.19	0.05
	(−0.57, −0.15)	(0.21, 0.94)	(−1.54, −0.47)	(−0.88, 0.46)
2008:Q1-2009:Q2	−0.91	−0.63	−3.01	−2.39
	(−1.46, −0.25)	(−1.76, −0.01)	(−4.21, −0.83)	(−4.71, −1.43)

Notes: Decomposition of Real GDP Growth into Demand and Supply Components during NBER Recessions. Individual GDP shocks decomposition are computed as $\sum_{i=1}^{i+N} \sigma_{gd} u_t^d$ (demand component) and $\sum_{i=1}^{i+N} \sigma_{gs} u_t^s$ (supply component), where i is the first time point of a recession which lasts for N quarters. Aggregate GDP shocks decompositions are computed by first computing $\sum_{i=1}^{i+N} g_t - E_{t-1} \sum_{i=1}^{i+N} g_t$ and $\sum_{i=1}^{i+N} \pi_t - E_{t-1} \sum_{i=1}^{i+N} \pi_t$ to identify the cumulative shocks to growth and inflation over the recession, where expectations are computed using the VAR. Supply and demand shocks are then computed using the inversion formula $\begin{bmatrix} \widehat{u}_t^s \\ \widehat{u}_t^d \end{bmatrix} = \widehat{M}^{-1} \begin{bmatrix} \widehat{u}_t^\pi \\ \widehat{u}_t^g \end{bmatrix}$. Units are real GDP, cumulative log difference x100. Figures in parentheses are 95 percent confidence intervals based on the bootstrapped covariance matrix for $\left[\widehat{A}, \widehat{M}\right]$.

Table 6: Decomposing the COVID recession

	GDP shock	demand component	supply component
2020:Q1	-1.63	-0.88	-0.75
	$(-1.75, -1.52)$	$(-1.20, -0.24)$	$(-1.42, -0.45)$
2020:Q2	-8.81	-4.57	-4.24
	$(-9.26, -8.57)$	$(-6.15, -1.17)$	$(-7.65, -2.70)$
2020:Q3	4.95	6.22	-1.27
	$(-0.83, 9.09)$	$(2.22, 7.73)$	$(-6.48, 6.11)$
2020:Q4	-1.17	-0.77	-0.40
	$(-2.18, 0.55)$	$(-1.34, -0.01)$	$(-1.90, 1, 24)$
2021:Q1	0.85	0.83	0.02
	$(0.76, 0.95)$	$(0.33, 0.98)$	$(-0.16, 0.51)$

Notes: Decomposition of Real GDP Growth into Demand and Supply Components during the COVID recession. Individual GDP shocks decompositions are computed as $\begin{bmatrix} \widehat{u}_t^s \\ \widehat{u}_t^d \end{bmatrix} = \widehat{M}^{-1} \begin{bmatrix} \widehat{u}_t^\pi \\ \widehat{u}_t^g \end{bmatrix}$, where $\widehat{u}_t^\pi$ and $\widehat{u}_t^g$ are the residuals for inflation and real GDP growth based on the VAR parameters in Table 2. Units are real GDP, log difference x100. Figures in parentheses are 95% confidence intervals based on the bootstrapped covariance matrix for $\widehat{M}$.

Figure 1: Reduced-form shocks

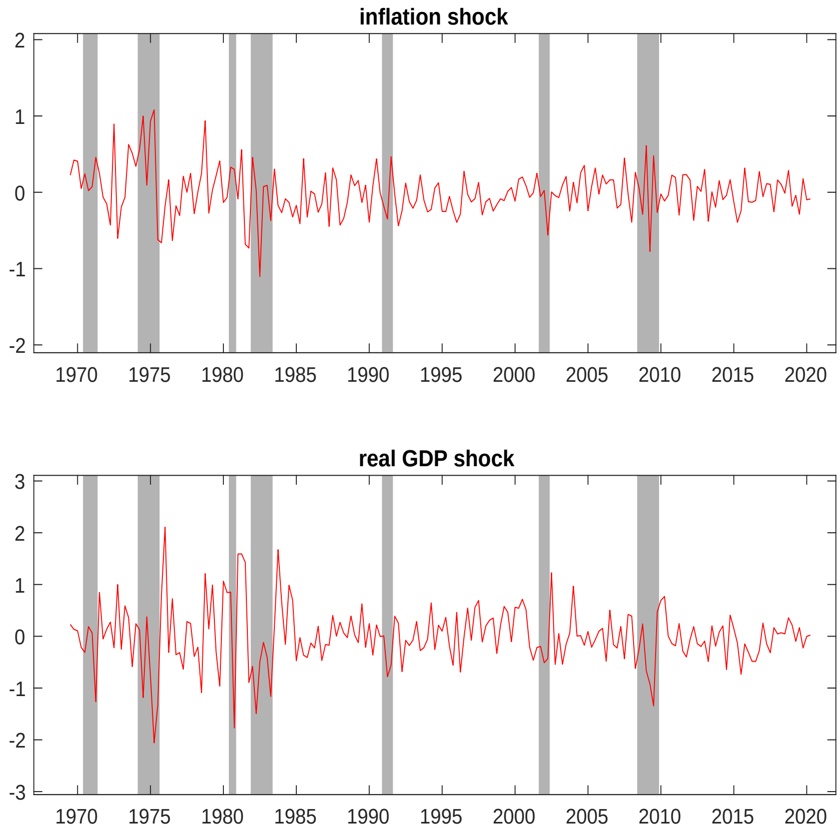

Notes: Data are quarterly 1968Q4-2020Q1. VAR setup is $Y_{t+1} = \mu + AY_t + u_t$. Variables include the log growth rate of real GDP, the log growth rate (multiplied by 400) of the GDP deflator, and one-quarter ahead forecasts for real GDP growth and inflation of the GDP price deflator, both from the Survey of Professional Forecasters. The top panel plots VAR residuals for inflation. The bottom panel plots residuals for real GDP growth. NBER recessions are shaded in gray.

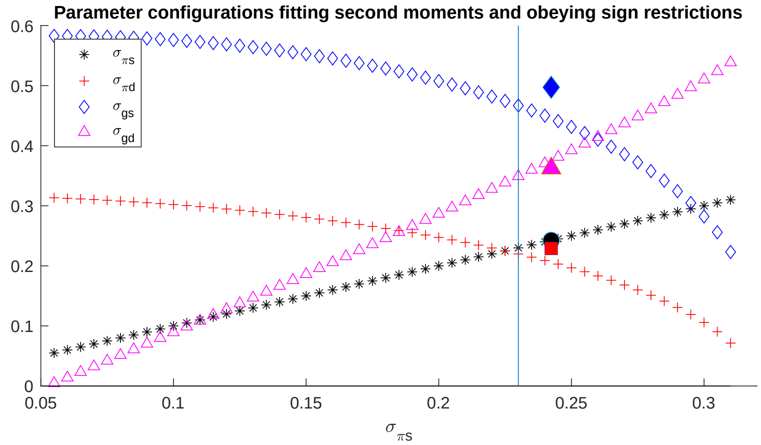

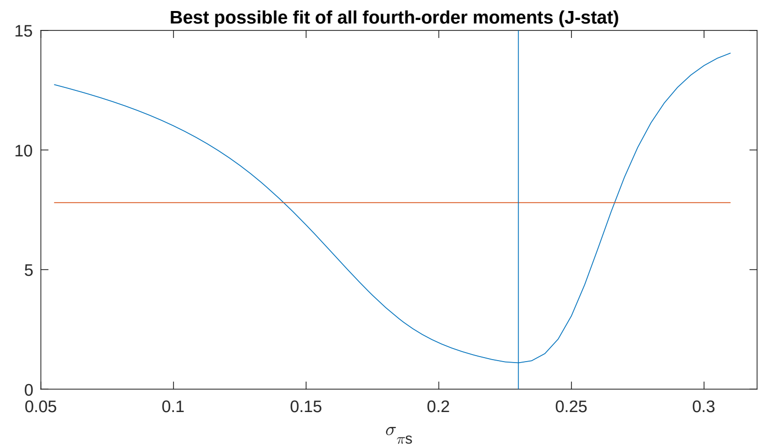

Notes: The lines in the top panel depict all possible configurations of the vector $[\sigma_{\pi s}, \sigma_{\pi d}, \sigma_{gs}, \sigma_{gd}]$ that are consistent with the sample covariance matrix reported in Table 3, as well as the sign restrictions. The bottom panel shows the J statistic results from the exercise: conditional on each configuration of $[\sigma_{\pi s}, \sigma_{\pi d}, \sigma_{gs}, \sigma_{gd}]$ in the top panel, $x\kappa_s$ and $x\kappa_d$, the excess kurtosis of supply and demand shocks, respectively, are estimated to fit the univariate kurtosis moments and the cross kurtosis moment reported in Table 3 as well as possible. The horizontal line in the bottom panels depicts the 5% rejection level for the $\chi^2(3)$. The large, solid, markers in the top panel represent the point estimates for the vector $[\sigma_{\pi s}, \sigma_{\pi d}, \sigma_{gs}, \sigma_{gd}]$ that result from the one-step estimation described in the text. The vertical line represents the minimum of the J-statistics in the bottom panel.

Figure 3: Recovered supply and demand shocks

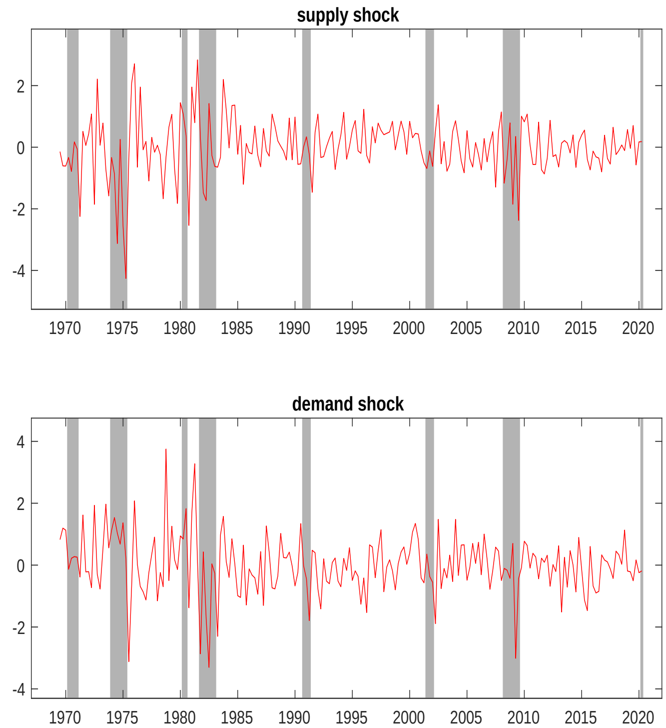

Notes: The top panel plots supply shocks and the bottom panel plots demand shocks estimated using the equation

$$\begin{bmatrix} \widehat{u}_t^s \\ \widehat{u}_t^d \end{bmatrix} = \widehat{M}^{-1} \begin{bmatrix} \widehat{u}_t^\pi \\ \widehat{u}_t^g \end{bmatrix}$$

where the reduced-form shocks u_t^π and u_t^g are recovered from the VAR and estimates of M are recovered from the one-step CMD estimation described in the text and point estimates reported in Table 3.