The Demand Origins of Business Cycles

Christian Matthes and Felipe Schwartzman*

February 15, 2019

Abstract

We use economic theory to rank the impact of structural shocks across sectors. This ranking helps us to identify the origins of U.S. business cycles. To do this, we introduce a Hierarchical Vector Auto-Regressive (Hi-VAR) model, encompassing aggregate and sectoral variables. We find that shocks whose impact originate in the "demand" side (monetary, household and government consumption) account for 2.4 times more of the variance of U.S. GDP growth at business cycle frequencies then identified shocks originating in the "supply" side (technology and energy). Furthermore, corporate financial shocks, which theory suggests propagate to large extent through demand channels, account for 1.4 times as much as those same supply shocks.

JEL Classification: C11, C50, E30
Keywords: Aggregate Shocks, Sectoral Data, Bayesian Analysis, Impulse Responses

1 Introduction

What drives business cycles? Macroeconomists have alternatively argued for demand factors such as monetary and fiscal expenditures shocks, and supply factors such as technological innovation and the cost of raw materials. So far, comprehensive decompositions of output fluctuations into the contributions of various shocks has only been obtained in tightly specified structural models. Those typically indicate a prominent role for supply shocks.[1]

We propose an approach to identify a variety of demand and supply shocks simultaneously, but within a flexible statistical framework. We identify shocks based on prior knowledge of their impact on different sectors. Thus, for example, an energy cost shock is identified with an aggregate shock that increases energy prices, and has a larger price and output impact on energy intensive sectors. Our analysis suggests a prominent role for shocks that manifest themselves in the demand side.

In order to implement this identification scheme, we introduce a large scale, flexible, and tractable econometric model, the Hierarchical Vector Auto-Regression (Hi-VAR). It allows us to analyze aggregate and sectoral time-series jointly, while allowing for rich internal sectoral dynamics. Aggregate shocks are captured by common factors in the innovations of the various series, and shock

*Federal Reserve Bank of Richmond. Contact information: christian.matthes@gmail.com, felipe.schwartzman@rich.frb.org. We would like to thank seminar participants at Fundacao Getulio Vargas as well as Ricardo Reis and Mark Watson for helpful comments. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of Richmond or the Federal Reserve System.

[1] Justiniano et al. (2010) attribute 75% of GDP fluctuations to "neutral" and investment specific TFP shocks.

identification is obtained by setting priors on factor loadings. The introduction of this particular econometric framework forms a separate methodological contribution of this paper.

In our baseline analysis, we identify six structural shocks: energy cost, technological progress, monetary, corporate finance, government consumption, and household demand. The first five shocks can be easily motivated with reference to an extensive literature. We take the household demand shock to encompass the set of shocks to household credit, wealth or expectations that have been most heavily emphasized following the 2007-09 recession, and which operate mainly through their effect on household consumption decisions. The identification scheme used for each of the shocks is done by analogy to the energy shock example given above, and mostly relies on input or demand intensity shares that can be read directly from input-output tables. The two exceptions are corporate credit shocks (tied to external financial dependence measures as in Rajan and Zingales (1998)) and monetary shocks (tied to sectoral price stickiness measures by Nakamura and Steinsson (2008)).

We check the reasonableness of our identification scheme by inspecting the resulting impulse response functions for aggregate variables. We find that those largely conform to theoretical priors and findings by studies focusing on single, well identified shocks.

Our results point to a prominent role for fluctuations originating in the demand side of the economy and in access to corporate credit. We find that demand side fluctuations (monetary, household demand and government consumption) account for a fraction of the variance of GDP growth at business cycle frequencies which is more than two times as large as the fraction accounted by identified supply (technology and energy) shocks. Of the three demand shocks, monetary policy features most prominently, followed by household demand and government consumption. Furthermore, corporate financial shocks, which theory implies propagate to a substantial degree through demand channels, account for 1.7 times as much as those same supply shocks.

The emphasis on the demand origins of business cycles complements results by Angeletos et al. (2018), who find that most business cycle fluctuations do not seem to be driven by supply shocks such as technological innovations. Those findings are in contrast to the results of prominent studies based on tightly specified DSGE models, which tend to emphasize supply factors.

A standard approach to identifying structural shocks is the use if sign restrictions, as poineered by Uhlig (2005), Faust (1998), and Canova and Nicolo (2002). The use of sign restrictions on sectoral responses to identify aggregate shocks processes has been analyzed by Amir-Ahmadi and Drautzburg (2017).[2] Relative to this previous work, our main innovation is to develop a method that allows us to investigate the role of several shocks simultaneously in a tractable manner in environments with large amounts of data. This paper thus also falls into a more general trend within macroeconomics of using cross-sectional data to inform on questions of relevance to macroeconomists (Holly and Petrella (2012) and Beraja et al. (2016)). We also add to an existing suite of time series models designed to incorporate large panels, including dynamic factor models Stock and Watson (2005)), factor augmented VARs (Bernanke et al. (2005), Boivin et al. (2009)) and global VARs (Chudik and Pesaran (2016), Holly and Petrella (2012)). Lastly, on a more technical note, we add to a literature that relies on Bayesian priors rather than hard identification restrictions (Kociecki (2010) and Baumeister and Hamilton (2015)), where our main contribution is to provide a method to add those restrictions in a maximally tractable manner.

The paper proceeds as follows: Section 2 describes the Hi-VAR model and the identification

[2] Also, Schwartzman (2014) and Fulford and Schwartzman (2015) use cross-sectional information to identify shocks. Whereas the first paper uses a structural small open economy model, the second paper leverages the cross-sectional impact of a shock identified from an historical narrative. The case for using information on the relative magnitude of the responses to shocks to help identify shocks has also been made by De Graeve and Karas (2010).

procedure. Section 3 lays out an analytically tractable multi-sector model with sticky prices and wages to provide the theoretical motivation for the identification assumption. Section 4 presents the results. Section 5 concludes.

2 Hierarchical VAR model: Identification and Estimation

We combine a VAR-type time series model for a vector of aggregate variables Y_t with autoregressive models for vectors of sectoral data X_t^i, where i indicates the sector. Aggregate and sectoral data interact in two ways: (i) via structural shocks that affect both types of data and (ii) via direct feedback from (lagged) aggregate data to sectoral data.[3] We use a Gaussian prior for the effects of the structural shocks on aggregate and sectoral data.[4] This procedure allows us to impose more prior information on the magnitudes of these effects compared to what would be feasible in the standard sign restriction approach.[5] By exploiting our specific model structure, we can efficiently estimate very large scale models. Also, because we directly estimate a structural VAR, our approach can handle set-identified, exactly identified, and over-identified environments - the differences just amount to to choosing different priors on the parameters governing the contemporaneous impact of structural shocks.

2.1 Modeling aggregate variables

We model aggregate variables as following a linear vector autoregressive process. A key difference from traditional VARs for aggregate data is that we break the tight link between forecast errors and structural shocks, thus allowing sectoral data to help identify structural shocks.

The aggregate variable vector Y_t (of dimension N by 1) is a function of its past values, structural shocks ε_t and other shocks w_t:

$$Y_t = \mu + \sum_{l=1}^{L} A_l Y_{t-l} + D\varepsilon_t + w_t \tag{1}$$

where ε_t is of dimension $S \times 1$, and Σ is an $N \times S$ matrix encoding the impact of the Gaussian structural shocks ε on the aggregate variables, and w_t is a $N \times 1$ vector of mean 0, non-structural Gaussian shocks with covariance matrix Ω. We further assume that $\varepsilon \sim N(0, I)$.[6] As will be clear later, we can allow for $S < N$, $S = N$, or $S > N$, whereas standard VAR analyses require $S \leq N$.

For later discussion, it is useful to note that the one-step ahead forecast error for the aggregate level is given by $D\varepsilon_t + w_t$, whereas a standard VAR model for the aggregate variables would assume

[3] In the appendix we discuss extensions that allow for more flexible feedback.

[4] We can do this because we directly estimate the impact of structural shocks rather than first estimate a reduced-form model and then infer the structural model afterwards, as is common in the VAR literature. In directly estimating a structural representation, we follow in the footsteps of, for example, Baumeister and Hamilton (2015) and Sims and Zha (1998), who directly estimate structural VARs.

[5] In the standard approach to impose sign restrictions, as outlined in Rubio-Ramirez et al. (2010), inequality restrictions are imposed on impulse responses in conjunction with a uniform (Haar) prior on the rotation matrices that map reduced form parameters to initial impulse responses. In the appendix we show how to incorporate strict inequality restrictions in our framework should a researcher be interested in those.

[6] The distributional assumptions are necessary because we ultimately want to carry out Bayesian inference, for which we need to build a likelihood function.

that any estimate of the structural shock is a linear combination of the aggregate one-step ahead forecast error.[7]

2.2 Modeling idiosyncratic variables

There are observations for I idiosyncratic units (such as industries, regions, or, in our specific application, sectors) with K variables (such as prices and quantities) each. The law of motion for the data from unit i, summarized in the K-dimensional vector X_t^i, is given by:

$$X_t^i = \mu^i + \sum_{l=1}^{L^X} B_l^i X_{t-l}^i + \sum_{l=1}^{L^Y} C_h^i Y_{t-l} + D^i \varepsilon_t + w_t^i \tag{2}$$

where now D^i is a $K \times S$ matrix encoding the impact of shocks ε on the idiosyncratic variables and the mean zero Gaussian vector w_t^i incorporates the impact of idiosyncratic (or non-structural) shocks on individual units. We denote the covariance matrix of w_t^i by Ω^i. We assume that w_t^i is independent across i and independent from w_t.

2.3 Interpreting our model

To gain further intuition, it will be useful to rewrite our model as follows: First define the vector of all observables

$$Z_t = [Y_t' \; X_t^{1'} \; X_t^{2'} \; \ldots \; X_t^{I'}]'$$

We can then recast our model in the following way:

$$Z_t = \mu^Z + \sum_{l=1}^{max(L^X, L^Y, L)} B_l^Z Z_{t-l} + \underbrace{D^Z \varepsilon_t + w_t^Z}_{u_t^Z} \tag{3}$$

where w_t^Z is a vector that stacks the non-structural shocks according to the ordering of observables in Z_t. Our model imposes restrictions on the matrices B_l by assuming that one sector's variables cannot directly respond to any other sector's lagged variables. Note that our one step ahead forecast error u_t^Z is iid and independent of all other right-hand side variables. We now provide a characterization of identification as it pertains to both the parameters of the model as well as the aggregate shocks. Because of the structure of the one-step ahead forecast error, we can identify μ^Z as well as the coefficient matrices B_l^Z. Even single equation-estimation via OLS would yield consistent estimates in our environment. Since we use even more information in our likelihood-based approach, the same identification results carry over. Focusing on u_t^Z, we can see that it follows a factor structure, where the common factors are the iid structural shocks ε_t. Note also that our assumptions on the correlation structure of w_t^Z limits the correlation of those idiosyncratic components across observables (these shocks can be correlated within sector and at the aggregate level, but not across the sets of equations outlined above). In fact, standard results on identification in Factor models apply (Bai and Ng (2008)). In fact, while the effects of individual structural shocks are not identified without additional assumption, the overall effect of *all* structural shocks

[7] This is true even if fewer than N shocks are identified, as is common in the literature on sign restrictions in VARs.

is identified, as we show below and later highlight in a Monte Carlo experiment. To identify ε_t, we need identification restrictions akin to those used in the structural VAR literature. To see this, define

$$u_t = D\varepsilon_t + w_t \tag{4}$$

$$u_t^i = D^i\varepsilon_t + w_t^i \; \forall i \tag{5}$$

For any conformable orthogonal matrix Q we can construct alternative models that feature the same first and second moments and thus the same Gaussian likelihood:

$$u_t = \underbrace{DQ^{-1}}_{\tilde{D}} \underbrace{Q\varepsilon_t}_{\tilde{\varepsilon}_t} + w_t \tag{6}$$

$$u_t^i = \underbrace{D^iQ^{-1}}_{\tilde{D}^i} \underbrace{Q\varepsilon_t}_{\tilde{\varepsilon}_t} + w_t^i \; \forall i \tag{7}$$

With the sign and magnitude restrictions in this paper we are not going to pin down a unique value of Q to get exact identification even though the overall impacts $D\varepsilon_t$ and $D^i\varepsilon_t \; \forall i$ are identified. However, even though we are in the realm of set identification, important recent work on the usefulness of restrictions of the kind we use (Amir-Ahmadi and Drautzburg (2017)) shows that they can be very informative. This is especially true for our setting, where we have many sectors on which we impose restrictions. What sets our approach apart from the previous literature on structural VARs is that (i) because of our model structure we can use substantially larger data sets than standard VAR applications can, (ii) for that same reason we can identify several shocks simultaneously, rather than one or two.[8]

Importantly, our approach is computationally very efficient. This is because, as we will show below, it relies solely on standard steps in Gibbs samplers (drawing from Normal and inverse-Wishart priors as described in Koop and Korobilis (2010) as well as using Gibbs sampling for linear and Gaussian state space models as in Carter and Kohn (1994)) that, in our specific case, are especially amenable to parallelization.[9] This implies that our approach can be very efficient even in applications that have a much larger scale than our application in this paper.[10] Finally, how do we interpret the shocks w_t? These are shocks that do not have a contemporaneous effect on sectoral data, while they do affect the aggregate data at time t. Furthermore, these innovations do not have an independent role in determining sectoral data beyond how they influence aggregate data. To safeguard ourselves against a scenario where identified w_t is estimated to be more important for determining aggregate data than it actually is in reality, we suggest to add additional shocks to the vector of structural innovations ε_t on which no identification restrictions are imposed. The additional 'structural' shocks will soak up any explanatory power that would otherwise falsely be

[8]By estimating the responses to structural shocks directly, we do not need to post-process reduced-form VAR estimates to obtain the structural representation that allows us to compute the effects of structural shocks. This is useful because the algorithms used to deliver the impulse responses after estimating a reduced-form model can be numerically inefficient because not all proposed candidate parameter vectors of the structural VAR satisfy the identification restrictions Rubio-Ramirez et al. (2010) or because the imposed restrictions are actually overidentifying as in Amir-Ahmadi and Drautzburg (2017).

[9]This parallelization argument does not hold, for example, in large scale VARs. And while certain aspects of Gibbs samplers for factor models might also be amenable to parallelization, these models do not directly emphasize the dynamics of all variables in sector in a transparent fashion.

[10]As a final note on the model, it might be helpful to note that one could interpret the lagged aggregate variables Y_{t-l} as additional, but observable, factors.

attributed to w_t. In our application, we add three of those shocks, but also show in the Appendix results with 10 additional shocks as a robustness check.

2.4 Setting Priors

An important step in our analysis is in the setting of priors. In contrast with traditional approaches, that achieve identification by setting hard constraints on the shock process, we follow Baumeister and Hamilton (2015) in using "soft" prior restrictions for identification. At the same time, by setting Gaussian priors on the direct impact of the structural shocks on variables and inverse Wishart priors on the variances, we can use a Gibbs sampler to estimate a large scale model with several identified shocks very efficiently.

2.4.1 Priors on D and D^i

The most important priors for our identification purposes are the ones we set on the impact matrices, D and D^i. We set priors for both parameters to be Gaussian. To set the average value for element k, s of D^i, we assume that it can decomposed as follows:

$$(E\left[D^i_{k,s}\right])^2 = \gamma^i_k \beta_{k,s} \alpha^i_{k,s} \quad , \tag{8}$$

$$(E\left[D^i_{k}\right])^2 = \sum_{s=1}^{S} (D^i_{k,s})^2 \quad . \tag{9}$$

where

1. $\alpha^i_{k,s}$: a measure of the *relative* impact of shock s on variable k *for sector i as compared to other sectors*. For example, this variable could encode the fact that a more energy intensive sector ought to be more sensitive to shocks pertaining to the price of energy than a less energy intensive one. This measure, which we derive from cross-sectoral data, is *not* comparable across shocks. To ensure that we are capturing the effect of the shock, we only use identifying information for those sectors that are in either extreme of the distribution of our indicator variables, remaining agnostic on the ones which lie in the middle. The specific indicators that we use to choose $\alpha^i_{k,s}$ are informed by the model presented in Section 3 and will be described in detail in Section 3.7.[11]

2. $\beta_{k,s}$: a measure of the *overall* impact of shock s on variable k *across all sectors*. For example, this variable could encode the notion that aggregate productivity shocks account for a larger fraction of the variance of quantities in all sectors as compared to markup shocks and vice versa for prices. We use an 'ignorance prior' and set this variable to $1/S$, where S is the number of structural shocks.

[11] To be precise, we carry out the calculation for all sectors and shocks in our model, but then only use that information for those sectors and shocks where the shock-specific measure of sectoral sensitivity is either in the bottom 10 percent or the top 90 percent. We do this because we only want to impose the most robust implication of our calculations as priors for our estimation. For those elements of $D^i_{k,s}$ that do not get restricted we use a loose prior with mean 0 and standard deviation 0.25, which is the same prior we use for the elements of D that are not restricted by theory. If there are missing values for the indicator variables for some sectors, we assume that the indicators for those sectors takes on the average value of the relevant indicator.

3. γ_k^i: a measure of the overall sectoral sensitivity to shocks. For example, this variable could encode the notion that consumption of durable goods is overall more sensitive to all shocks than consumption of nondurables. Given $\alpha_{k,s}^i$ and $\beta_{k,s}$, we can back out this variable if we have values for $(D_k^i)^2$. We obtain those by estimating the model in a training sample from our estimation with an agnostic prior. We do not need to impose any identifying restrictions on the structural shocks for that training sample step because we are only interested in estimating how important those shocks are for fluctuations of different variables together. The factor structure of the shocks allows us to do that even if we cannot disentangle the individual shocks. We also only impose very loose prior on the covariance matrix of the non-structural shocks.

The procedure above allows us to set a magnitude for the prior mean $E\left[D_{k,s}^i\right]$. We use a priori information on the sectoral impact of shocks to set the sign.

To set the prior mean for the impact of shocks on aggregate $E\left[D_{ks}\right]$, we only use minimal assumptions: a monetary policy shock is more likely to raise the nominal rate than to lower it, a household demand shock tends to increase consumption, a government spending shock tends to increase government spending, a technology shock tends to increase measured TFP, a credit shock shock tends to increase our measure of the spread, and finally an energy shock tends to increase energy price inflation. Those prior assumptions, together with the priors regarding the impact of the shocks on sectoral variables are summarized in table 2 For the impact of those shocks on other variables, we impose an agnostic prior with mean zero and variance 0.25. Given those assumptions, the training sample pins down the prior mean impact of the shocks on the aggregates.

In order to obtain a prior standard deviation for $D_{k,s}^i$, we choose the prior standard deviation to be $k \times abs\left(E\left[D_{k,s}^i\right]\right)$, where we set k to be 0.1.

At this point it is useful to step back. What we have solved for here is the *prior mean* of $D_{k,s}^i$. Let's make this explicit and call the (square root of the) solution to the system of equations above $E\left[D_{k,s}^i\right]$. This means that we do not necessarily force the sign restrictions to hold with certainty. Since prior on $D_{k,j}^i$ is normal there will always be some probability that the sign of the posterior mean will be different from that of the prior mean. This also highlights the role of the prior variance.[12]

2.4.2 Setting priors for Ω and Ω^i

To use the Gibbs sampler, we use inverse-Wishart priors for the covariance matrices of the reduced form shocks at the aggregate and sectoral levels. As is well known, this imposes some restrictions on what prior beliefs we can impose on our model. One is that the variances are bounded away from 0 (really not much of a problem in our case), while the main problem is that there is no truly uninformative prior (as we increase the variance, we also have to at some point increase the prior mean since variances are bounded below by 0).

Since the priors for these covariances turn out to be potentially influential in finite samples, we will use the results of the estimation with the agnostic prior to set this prior.

[12] In order to preserve the magnitude of $(D_k^i)^2$ when we only use certain sectors to inform our prior, we can rescale $D_{k,s}^i$ for the relevant i, k, and s indices to match the original magnitude of $(D_k^i)^2$. In our specific application this did not make a material difference in the results.

Prior on Ω To set the prior for Ω, we use results from our agnostic prior estimation. We set the prior mean to the estimated posterior mean of Ω and use as degrees of freedom the size of our overall sample.

Prior on Ω^i For Ω^i, we follow the same strategy as for its aggregate counterpart Ω.

2.5 Sampling Strategy

As mentioned before, we exploit the Gibbs sampler throughout by imposing independent Normal-inverse Wishart priors.

2.5.1 Drawing ε_t given all other parameters

We assume Gaussian innovations throughout for tractability. If we use a variant of equations (1) and (11), it is straightforward to see that, conditional on A_l, B_l, C_l, Σ, D and D^i, ε_t can be drawn via exploiting the Kalman filter (simply put all known quantities on the left-hand-side: all that remain on the right hand side are the ε terms, w^i and w), based on Carter and Kohn (1994). To make this step more numerically efficient, we follow Durbin and Koopman (2012) and collapse the large vector of observables in a vector with the same dimension as the structural shocks. As discussed by Durbin and Koopman (2012), this can be done without loss of information.

2.5.2 Drawing other parameters given ε

Since we condition on ε at this stage, drawing all other parameters amounts to drawing from Gaussian and inverse Wishart posteriors. For the aggregate equations, we use a Minnesota-type prior followingKoop and Korobilis (2010). Such a prior is useful because it avoids overfitting. Since there are only two variables per sector, overfitting is less of an issue, so we use priors centered at 0 with a standard deviation of .5 there. One helpful insight here is that conditional on ε, all other blocks can be run in parallel. This means that our approach can be scaled up easily. This is especially useful for extensions where the researcher might want to depart from the Normal-inverse Wishart prior used here.

2.6 Comparison with other approaches

Our model differs from other approaches that try to model large panels of time series by explicitly modeling a distinction between time series at the aggregate and idiosyncratic levels. FAVARs (Bernanke et al. (2005)) do feature this distinction, but do not explicitly model dynamics at the sectoral level. Furthermore, identifying structural shocks in both factor models and FAVARs requires imposing identifying assumptions for both the unobserved factors and the structural shocks which, in turn, identify the factors. While our model does have a factor structure, we only require identifying assumptions for the structural shocks since all other factors are observable. Another modeling approach that touches on issues similar to ours are Global VARs (Chudik and Pesaran (2016)). Those do not break the tight link between aggregate shocks and one-step ahead forecast errors at the aggregate level and require restrictions on how shocks propagate between idiosyncratic variables. At this point the reader might wonder why breaking the tight link between one-step ahead forecast errors and structural shocks implied by standard VARs is useful. There are two distinct

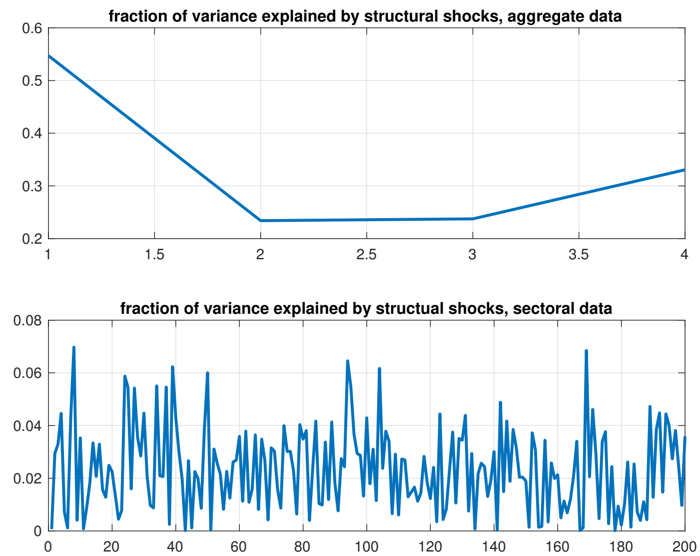

Figure 1: Fraction of variance explained by structural shocks, Monte Carlo experiment

reasons: (i) this allows sectoral data and aggregate data to *jointly* identify structural shocks and (ii) it does not necessarily force structural shocks to explain large fractions of the variances of our observables if the data not call for structural shocks to be important.[13]

2.7 A Monte Carlo Experiment

This section describes the results of a Monte Carlo experiment that is meant to highlight that the overall impact of structural shocks in our environment is identified *independently of identifying restrictions for any specific structural shock*. All variables in this example are stationary, even though this is not necessary for our method. The aggregate level consists of 4 variables, whereas each sector (of which we have 100) consists of 2 variables. There are two structural shocks (elements of ε_t). The additional, non-structural, shocks are correlated within units (sectors or the aggregate level). We assume that most of the variance of the one-step ahead forecast errors is in fact due to these additional shocks - Figure 1 displays the fraction of the one-step ahead forecast error due to structural shocks for the variables at the aggregate and sectoral levels. These numbers are meant to convey that this is in fact a hard inference problem - most of the variation in the simulated data is not due to structural shocks. We simulate 130 datapoints and assume a lag length of 1 in all specifications (which is the correct specification). We use the agnostic prior for this estimation.

The Gaussian priors for all coefficients are centered at 0 with a variance of 10 and are thus loose given the magnitude of the parameters used in the estimation. A loose Wishart prior is used for the covariance of the residual error terms.

Even without *any* specific identifying restrictions on structural shocks, the next two figures show

[13]Our model does not preclude structural shocks from being the main drivers of business cycles a priori: The estimated variances of the non-structural shocks could be very small.

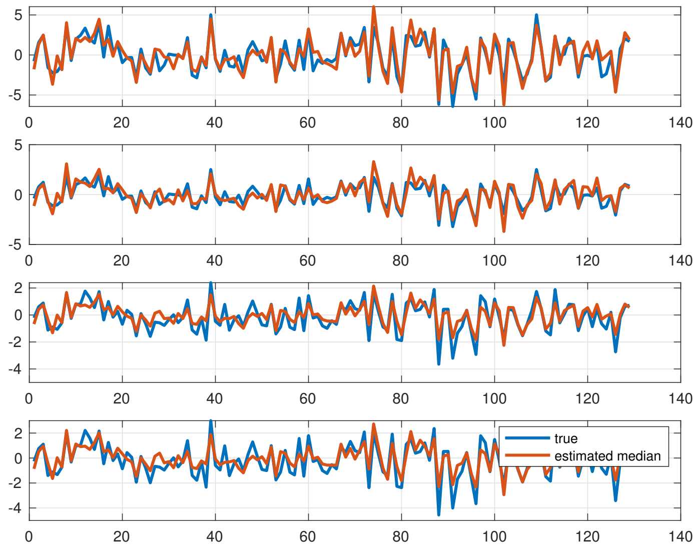

Figure 2: True effect and estimated median effect, Monte Carlo experiment

that in finite samples, our approach is able to correctly predict the overall effect of structural shocks. Figure 2 plots the true and estimated (median)effects of structural shocks on aggregate variables. As outlined in our discussion of identification, these results highlight that while the individual effects of structural shocks can not be estimated without identifying assumption in our framework, the overall effect of all structural disturbances is well identified.

3 A Tractable Multi-Sector Model with Nominal Rigidites

We now lay out a tractable, multi-sector model with nominal rigidities to motivate the shock identification. Nominal rigidities allow for a non-trivial "aggregate demand" channel. Since our main focus is in the cross-sectional differences between industries rather than their individual dynamics, we lay out a static multi-sector economy. We also only impose identifying restrictions on the impact responses. The model shares many elements with the framework developed in Pasten et al. (2018), while also allowing for nominal wage stickiness and for several aggregate shocks.

3.1 Households

There is a representative household with Cobb-Douglas preferences over the various goods, with share-parameter α_i for a good of industry i.

$$U = \prod C_i^{\alpha_i},$$

where $\sum_i \alpha_i = 1$. The household chooses its the amount it consumes of good i, C_i, to maximize its utility subject to the budget constraint

$$\sum_i P_i C_i + T = WL + \Pi,$$

where T is a lump-sum tax levied by the government to finance its consumption, W is the wage rate, Π are profits rebated from firms and $L < 1$ is employment to be determined in equilibrium. Households supply one unit of labor inelastically, but nominal wages are rigid so that labor is rationed.

Optimal household consumption choice satisfies:

$$P_i C_i = \alpha_i^C P^C C,$$

for $P^C \equiv \prod_i \left(\frac{P_i}{\alpha_i}\right)^{\alpha_i}$ and $C \equiv \prod_i (C_i)^{\alpha_i}$

3.2 Fiscal Authority

The fiscal authority minimizes the cost of consuming a given some aggregate government consumption G,

$$\min P_i G_i$$
$$s.t. : \prod (G_i)^{\alpha_i^G} = G,$$

where G is exogenously determined and α_i^G are the shares. The optimality condition for the government is:

$$G_i = \alpha_i^G \frac{P_G}{P_i} G$$

where

$$P_G = \prod_i \left(\frac{G_i}{\alpha_i^G}\right)^{\alpha_i^G}.$$

3.3 Firms

There are I sectors, indexed $i \in \{1, ..., I\}$. Within each sector there is a continuum of varieties of intermediate products indexed $v \in [0, 1]$. Those varieties are purchased by final goods producers that bundle them into the I goods according to a CES aggregator:

$$Y_i = \left[\int_0^1 Y_i(v)^{\frac{\theta-1}{\theta}} dv\right]^{\frac{\theta}{\theta-1}}$$

The demand for final good producer in sector i for intermediate input of variety v is

$$Y_i(v) = \left(\frac{P_i(v)}{P_i}\right)^{-\theta} Y_i$$

where

$$P_i = \left[\int P_i(v)^{1-\theta} dv\right]^{\frac{1}{1-\theta}}$$

For each variety, production takes place with a Cobb-Douglas production function:

$$Y_i(v) = e^{\epsilon_i} \prod_j (X_{ji}(v))^{\gamma_{ji}} \times (L_i(v))^{\lambda_i} (K_i(v))^{\chi},$$

where $M_{ji}(v)$ is the quantity of materials produced in sector j used in sector i for variety v, $L_i(v)$ is labor, $K_i(v)$ is sector-specific capital, and ϵ_i is a sector-specific exogenous productivity shock. The share parameter for good j used in sector i is γ_{ji} and the share of intermediate inputs in production is γ_i. We assume that $\sum_j \gamma_{ji} + \lambda_i + \chi = 1$, so that firms in the industry face constant returns to scale.

Producers of varieties are monopolists. Letting $\mathbf{s} = \left\{m^G, m^C, m^Y, \{k_i\}_{i=1}^I, \{\epsilon_i\}_{i=1}^I\right\}$ denote the state of the economy, they take the wage rate, prices, and household demand as given and choose their inputs to maximize profits. Firms differ on the information set available to them regarding prices and the demand for their intermediate input.

$$\max_{M_{ji}} E\left[P_i(v)Y_i(v,\mathbf{s}) - \sum_j P_j(\mathbf{s})X_{ji}(v,\mathbf{s}) - w(\mathbf{s})L_i(v,\mathbf{s}) - r_i(\mathbf{s})K_i(v,\mathbf{s})|\mathcal{I}_i(v)\right]$$

$$s.t. :Y_i(v,\mathbf{s}) = \left(\frac{P_i(v)}{P_i(\mathbf{s})}\right)^{-\theta} Y_i(\mathbf{s})$$

$$Y_i(v,\mathbf{s}) = e^{\epsilon_i} \prod_j (X_{ji}(v,\mathbf{s}))^{\gamma_{ji}} (L_i(v,\mathbf{s}))^{\lambda_i} (K_i(v,\mathbf{s}))^{\chi}$$

where $\mathcal{I}_i(v)$ is the information set for variety v in sector i. For a fraction ϕ_i of variety producers in sector i ($v \in [0,\phi_i]$) the information set does not includes the realized vector of shocks $\mathbf{s}$. For the remainder, the information set does includes it. Yet, firms commit to producing as much as necessary to satisfy demand at the prices that they choose.

Given cost-minimization, marginal cost is

$$\mathrm{mc}_i(\mathbf{s}) = e^{-\epsilon_i} \prod_j \left(\frac{P_j(\mathbf{s})}{\gamma_{ji}}\right)^{\gamma_{ji}} \left(\frac{w(\mathbf{s})}{\lambda_i}\right)^{\lambda_i} \left(\frac{\mathbf{r}(\mathbf{s})}{\chi}\right)^{\chi}$$

Firms with full information set

$$P_i(v,\mathbf{s}) = \frac{\theta}{\theta-1}\mathrm{mc}_i(s)$$

Firms without full information set

$$P_i(v) = \frac{\theta}{\theta-1} E\left[\frac{P_i(\mathbf{s})^{\theta} Y_i(\mathbf{s})}{E\left[P_i(\mathbf{s})^{\theta} Y_i(\mathbf{s})\right]} mc_i(\mathbf{s})\right]$$

We thus have that the price index for sector i is

$$P_i(\mathbf{s}) = \left[\phi_i \frac{\theta}{\theta-1} E\left[\frac{P_i(\mathbf{s})^\theta Y_i(\mathbf{s})}{E\left[P_i(\mathbf{s})^\theta Y_i(\mathbf{s})\right]} mc_i(\mathbf{s})\right] + (1-\phi_i)\left(\frac{\theta}{\theta-1}\mathrm{mc}_i(\mathbf{s})\right)^{1-\theta}\right]^{\frac{1}{1-\theta}}$$

Given that all firms in a sector have the same marginal cost, we can write the average markup as

$$\mu_i = \frac{P_i(\mathbf{s})}{mc_i(\mathbf{s})} = \left[\phi_i \frac{\theta}{\theta-1} E\left[\frac{P_i(\mathbf{s})^\theta Y_i(\mathbf{s})}{E\left[P_i(\mathbf{s})^\theta Y_i(\mathbf{s})\right]} mc_i(\mathbf{s})\right] \frac{1}{mc_i(\mathbf{s})} + (1-\phi_i)\left(\frac{\theta}{\theta-1}\right)^{1-\theta}\right]^{\frac{1}{1-\theta}}$$

3.4 Market Clearing

Market clearing for each sector i, requires that all output is used either as materials, for household consumption or for government consumption:

$$Y_i = \sum_j X_{ij} + C_i + G_i$$

Also, there is a fixed stock of capital $\bar{K}_i$ for each sector. Market clearing in capital markets thus requires that the demand for capital in sector i equals supply:

$$K_i = \bar{K}_i$$

The resource constraint in the labor market is

$$\sum_i L_i \leq 1$$

With sticky wages the inequality need not hold. We assume that wages are stuck at a level high enough that it doesn't bind. Labor rationing thus implies that

$$L = \sum_i L_i$$

3.5 Shocks

There are two nominal quantities set exogenously: nominal private consumption and nominal government consumption. Specifically, we assume that

$$P^C C = M^C M^Y$$
$$P^G G = M^G M^Y$$

so that nominal private and government consumptions can be affected either by an exogenous component which is specific to each type of final expenditure M^C or M^G, or by a common component M^Y.

We also allow for industry level productivity shocks ϵ_i. We assume that $\epsilon_i = \sum_{r=1}^{R} \lambda_{ir}\epsilon_r + \hat{\epsilon}_i$, where ϵ_r are aggregate shocks, F_i captures the sensitivity of various sectors to that shock and $\hat{\epsilon}_i$ is a sector-specific shock. In our application, we will allow ϵ_r to incorporate shocks to technology but also financial shocks or shocks to the price of imported inputs.

3.6 Reduced log-linearized system

After log-linearizing the model and rearranging, the model can be reduced to:

$$
\begin{aligned}
p_i &= \frac{1-\phi_i}{1-\chi}\chi\left[-\epsilon_i + \sum_j \gamma_{ji}p_j + \chi(y_i - \bar{k}_i)\right] \\
p_i + y_i &= \sum_j \gamma_{ij}\frac{Y_j}{Y_i}\left(y_j + \frac{1}{1-\phi_j}p_j\right) + \frac{C_i}{Y_i}\left(m^C + m^Y\right) + \frac{G_i}{Y_i}\left(m^G + m^Y\right) \\
c_i + p_i &= m^C + m^Y
\end{aligned}
$$

The first set of equations are "sectoral supply" equations, relating marginal production cost to prices. The second set of equations are "sectoral demand"'equations, relating nominal expenditures to sectoral prices. The last set of equations link nominal consumption expenditures and exogenous demand shocks.

The system has the form

$$
Z = AZ + b = A^N Z + \sum_{n=0}^{N-1} A^n b
$$

with Z including prices and quantities in all sectors, b including the direct impact of all exogenous shocks and A including the indirect impact of shocks through linkages.

The direct impact of shocks is

$$
\begin{aligned}
p_i^{\text{Direct}} &= \Phi_i\chi\left[\frac{C_i}{Y_i}m^C + \frac{G_i}{Y_i}m^G + m^Y\right] - \Phi_i\left(\epsilon_i + \chi\bar{k}_i\right) \\
y_i^{\text{Direct}} &= (1-\chi\Phi_i)\left[\frac{C_i}{Y_i}m^C + \frac{G_i}{Y_i}m^G + m^Y\right] + \Phi_i\left(\epsilon_i + \chi\bar{k}_i\right)
\end{aligned}
$$

where

$$
\Phi_i \equiv \frac{1-\phi_i}{\chi(1-\phi_i) + 1 - \chi}
$$

is inversely related to ϕ_i. Indirect effects are

$$
\begin{aligned}
p_i^{\text{Indirect}} &= \Phi_i \sum_j \left(\chi\frac{f_{ij}}{1-\phi_j} + b_{ji}\right)p_j + \chi\Phi_i\sum_j f_{ij}y_j \\
y_i^{\text{Indirect}} &= (1-\chi\Phi_i)\sum_j f_{ij}y_j + \sum_j\left[\frac{1-\chi\Phi_i}{1-\phi_j}f_{ij} - \Phi_i b_{ji}\right]p_j
\end{aligned}
$$

where $f_{ij} = \gamma_{ij}\frac{Y_j}{Y_i}$ capture forward linkages and $b_{ij} = \gamma_{ij}$ captures backward linkages. Note that if $f_{ij} \simeq 0$, the impact of the shock is similar to that of a productivity shock. We take this to be a reasonable approximation for the impact of energy shocks.

3.7 Priors on sectoral impact of aggregate shocks

Given the model above, we use the following priors for the impact of structural shocks (signs for the prior mean in parentheses):

- Household demand: $\alpha^s_{i,k} = \frac{C_i}{Y_i}$ (+ for quantities and prices), obtained from Use tables published by the BEA

- Government demand: $\alpha^s_{i,k} = \frac{G_i}{Y_i}$ (- for quantities, + for prices), obtained from Use tables published by the BEA

- Credit: $\alpha^s_{i,k} = \partial\bar{k}_i/\partial credit$ proxied by dependence on external finance (- for quantities, + for prices) calculated from COMPUSTAT data.

- Energy: $\alpha^s_{i,k} = \gamma_{\text{energy},i}$ (- for quantities, + for prices) , obtained from Use tables published by the BEA

- Monetary: $\alpha^s_{i,k} = \phi_i$ for quantities and $1/\phi_i$ for prices, proxied by average price duration, obtained from Nakamura and Steinsson 2008.

- technology shocks: $\gamma^s_{i,k} = \partial\epsilon_i/\partial technology$ proxied by R&D intensity (+ for quantities, - for prices). Here we us the ratio of intermediate goods from tech industries to total intermediate goods following Heckler (2005)

Those prior assumptions, together with the priors regarding the impact of the shocks on aggregate variables are summarized in table 2

4 Data and Results

4.1 Data

We use 8 aggregate US time series (in year-over-year growth rates where applicable): real GDP growth, CPI inflation, the effective Federal Funds rate, growth rate in real government spending, real PCE consumption growth, Moody's Seasoned Baa Corporate Bond Yield Relative to Yield on 10-Year Treasury Constant Maturity, Fernald's utility adjusted TFP (Fernald (2014)), and energy inflation based on the relevant producer price index. We use data from the first quarter of 1961 to the first quarter of 2018.
For the sectoral data we use the growth rate of real activity as measured by the sectoral PCE

and sectoral inflation as measured by the associated price. In terms of specification, we use 6 lags throughout, except for the lagged aggregate variables in the sectoral equations, where we only use one lag. We allow for 6 aggregate shocks: monetary policy, government spending, financial, energy, technology, household demand. We also allow for 3 unidentified sources of sectoral variation, which we include to allow for the possibility that we missed some important aggregate structural shocks. We plot the data we use in the appendix.

Table 1 below summarizes the priors on the different parameters, whereas table 2 summarizes the aggregate variable and sectoral indicators used to pin down the prior means for the impact matrices.

Table 1: Summary of prior distributions

Parameters	Prior Density	Prior Parameters
μ, A_l	Normal	Minnesota prior as in Koop & Korobilis
Ω	Inverse Wishart	mean set via training sample, degrees of freedom set to sample size
D, constrained elements	Normal	mean and standard deviation set by solving system of equations
D, unconstrained elements	Normal	mean 0, standard deviation 0.25
μ^i,B^i_l, C^i_h	Normal	each element has mean 0, standard deviation 0.5
Ω_i	Inverse Wishart	prior mean set via training sample, degrees of freedom set to 15
D^i, constrained elements	Normal	mean and standard deviation set by solving system of equations
D^i, unconstrained elements	Normal	mean 0, standard deviation 0.25

Table 2: Summary of prior means

Shock	aggregate impact	Sectoral impact index
Technology	Fernald TFP	high-tech content
Credit	credit spread	financially dependence
Household	household consumption	consumption oriented
Government	government consumption	government oriented
Monetary	fed funds rate	price stickiness
Energy	energy price index	energy intensity

4.2 Impulse Response Functions

We now show the impulse response functions to different shocks. This provides a check on our identification procedure, in that it allows us to evaluate whether responses to identified shocks conform to theory or prior findings based on other identification schemes. To economize on space, we show the 5th percentile, median and 95th percentile of the impulse responses to a one-standard deviation shock for the monetary and household demand shocks in figures 3 and 4. For those shocks we will also highlight how the different components of our identification strategy (on aggregate and sectoral data) interact. For all other shocks, figure 5 shows the median IRFs. The full set of impulse response plots with error bands and the associated figures analyzing the role of our identification assumptions for all other shocks can be found in the appendix. The plots look largely as expected: technology shocks increase consumption and GDP without much of an effect on consumption,

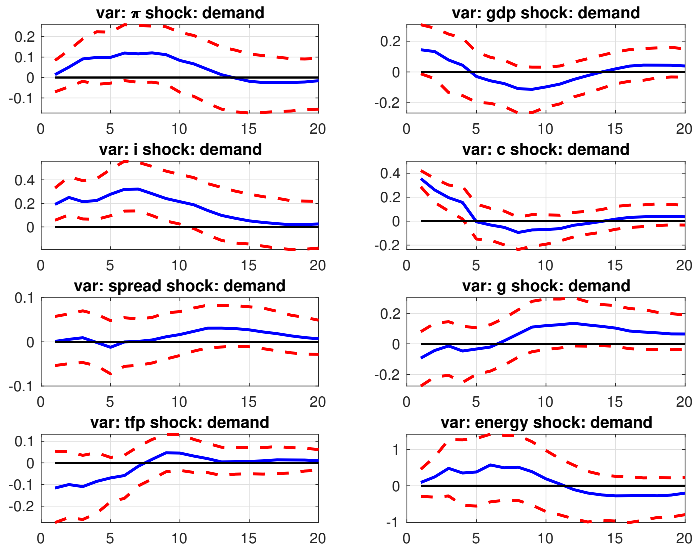

Figure 3: Responses to Household Demand Shock

while credit shocks depress those. Household demand shocks increase interest rate and GDP at first, the impact on output quickly reverts and the point estimate becomes somewhat negative (Figure 3). Government consumption shocks increase interest rates and inflation, and while they boost consumption at first, they crowd it out later. Energy shocks increase interest rates and inflation and depress GDP and consumption. Monetary policy shocks depress inflation, GDP and consumption (Figure 4).

We also show how incorporating the sectoral data helps with identification. For example, figure 6 shows that, relative to a specification where the shock is identified only from its impact on aggregate consumption, the impulse response functions for the household demand shock becomes much more tightly estimated once we incorporate priors on the sectoral responses. It is those tighter posteriors that make clear the impact of those shocks on inflation and interest rates. Also, figure 7 does the same analysis for monetary policy shocks. Compared to a shock identified solely from its impact on the nominal interest rate, the sectoral identification shows the deflationary impact much more clearly.[14]

[14]These figures plot 16th and 84th percentiles.

Figure 4: Responses to Monetary Shock

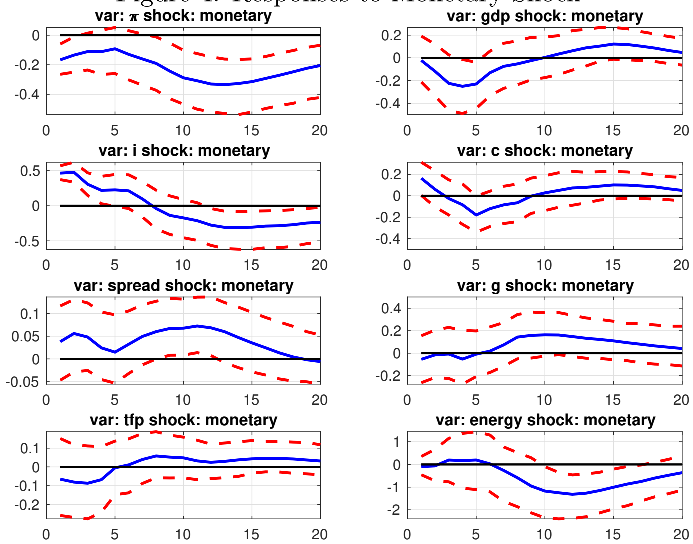

Figure 5: Median Responses to all other Shocks

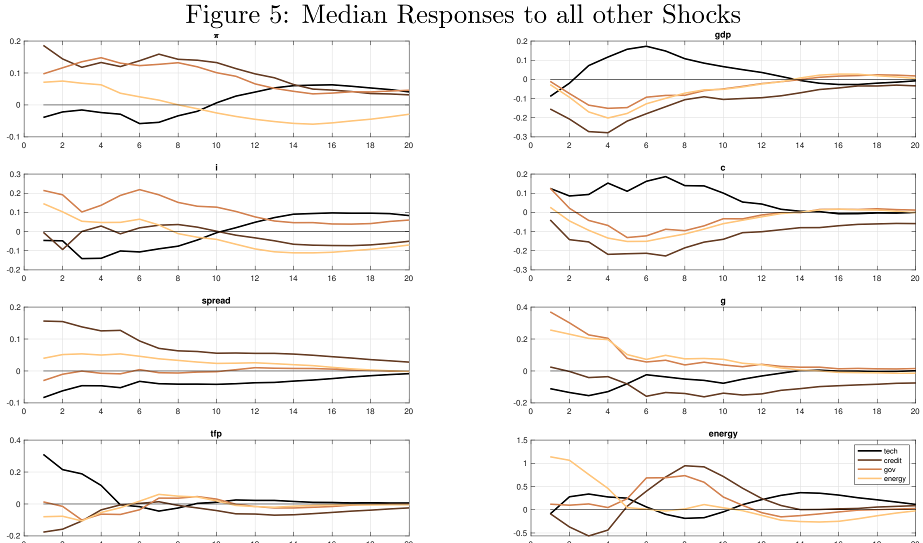

Figure 6: Responses to Household Demand Shocks: Comparison of Identification Schemes

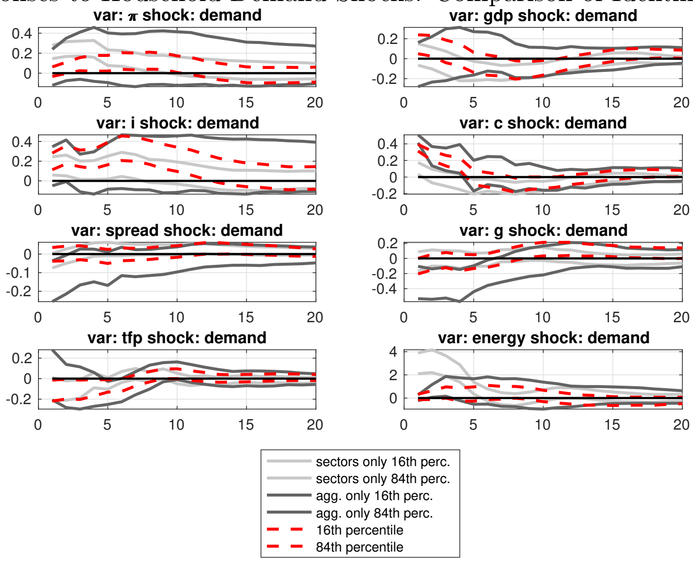

Figure 7: Responses to Monetary Shocks: Comparison of Identification Schemes

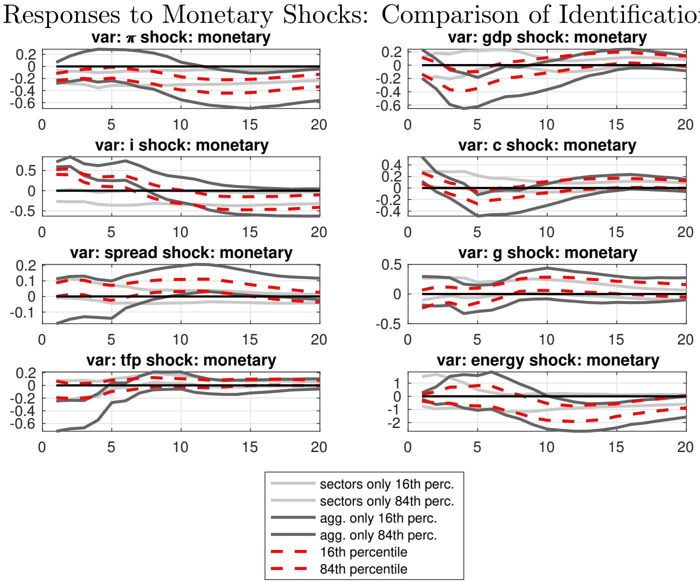

4.3 Analysis of Variance

In this section we show how much our identified structural shocks explain as a fraction of the variance explained by all structural shocks. The results are presented in table 3 below. To obtain the numbers in the table, we decompose for each variable the fraction of the variance of its innovation that is tied to aggregate shocks into different components. The numbers refer to average variances for forecast errors 6 to 32 quarters ahead. The six identified shocks account for more than 80% of overall variance explained by the structural shocks ε, with the one exception being Fernald's TFP series. The table shows that monetary shocks play a prominent role not only in explaining nominal interest rates and inflation (as one would expect) but also GDP, consumption and energy prices. The other shock with a prominent role is to corporate credit, accounting for a large part of the variance of GDP and consumption. If we count household household consumption, government consumption and monetary policy as "demand" shocks and energy and technology as "supply" shocks, we find that demand shocks account for 2.4 more of GDP variation then supply shocks.

	tech	credit	household	gov	energy	monetary	others
π	5.6	14.8	7.0	9.0	5.7	42.8	15.1
gdp	10.0	23.5	14.8	8.0	7.7	20.3	15.7
i	6.7	10.8	13.1	10.3	7.9	36.4	14.8
c	9.2	19.7	26.4	8.3	5.9	19.0	11.5
spread	13.1	37.4	7.1	6.1	7.1	19.1	10.1
g	7.1	18.6	12.3	17.9	12.0	16.3	15.8
tfp	21.7	16.1	7.7	4.7	6.0	11.7	32.1
energy	7.3	15.5	9.3	8.5	16.3	27.1	16

Table 3: Mean of variance decomposition across business cycle frequencies and posterior draws

4.4 Model Fit and the Interplay between Sectoral and Aggregate Data

Our model is highly restrictive in that any correlation between sectors as well as between sectors and aggregate variables has to come through either the structural shocks ε_t or lagged aggregate variables. The reader might a-priori wonder if this leads to substantial misspecification, which in turn would cast doubt on our identification strategy that is based on sectoral data.

To address this possible concern, we first compute the correlation between aggregate consumption growth and consumption growth at the sectoral level that appear in our dataset as well as the corresponding correlations for aggregate and sectoral inflation. We then draw 1000 parameter values from the posterior, simulate data of the same length as our dataset for each set of parameters (after discarding 1000 burn-in observations) and compute the same correlations for our simulated data. This gives us the posterior distribution of the correlations we are interested in. We are thus carrying out a posterior predictive check as advocated for by Rubin (1984) and further discussed by Gelman et al. (2013) and Geweke (2005), for example.

Figure 8 plots the correlations from the data (black) as well as the median (red) and the 5th and 95th percentiles (blue) of the posterior distribution. We sorted the correlations from the actual data by size (starting with the largest correlation) to make the figure easier to interpret. We order the sectors the same way for the simulated data. As can be seen from from figure 8, our model is able to replicate the correlation patterns between aggregate and sectoral data. An inquisitive

Figure 8: Correlations between sectoral and aggregate data, sectors on x-axis

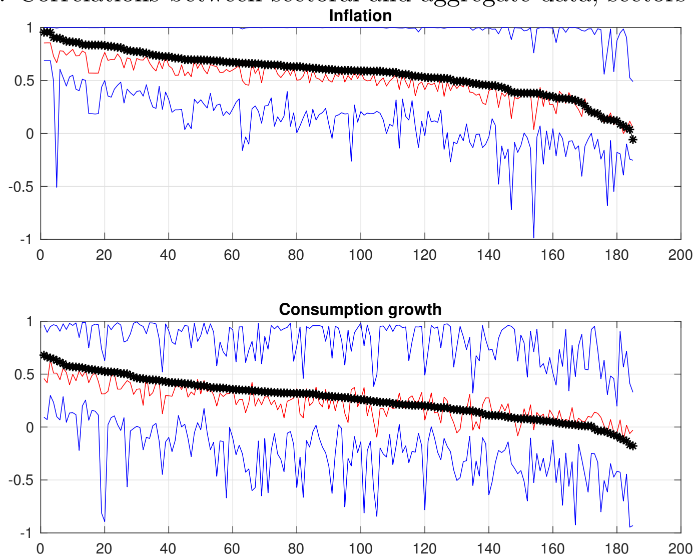

reader might ask for a more stringent test, namely a check of the correlation of variables *across* sectors rather than between any sector and the corresponding aggregate variable. We show the results for this posterior predictive check in figure 9. The figure looks noisier just because there are many more datapoints (pairwise correlations between the 185 sectors in our sample), but the main pattern remains, our model is able to replicate the broad correlation patterns. Our model misses at the very tail ends of the spectrum of correlations (more so for inflation than for consumption growth), but given that our model is tightly parametrized and parsimonious, we think of these results as very encouraging.

Figure 9: Correlations between sectoral data, sectors on x-axis

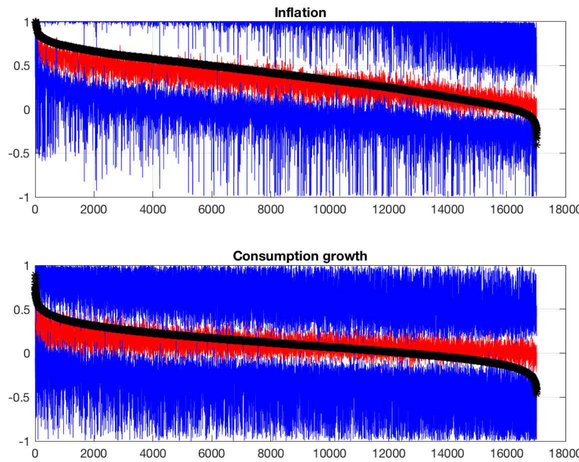

5 Conclusion

We lay out a new methodology to identify the effect of aggregate shocks and their role in driving aggregate fluctuations. The hierarchical vector auto-regressive model allows us to isolate innovations to aggregate and sectoral variables. We can then use those innovations to extract common factors that isolate the role of aggregate shocks in driving fluctuations. Those factors are identified through priors on their differential impact on different sectoral prices and quantities combined with a priori identification of those shocks with a single aggregate variable.

This identification procedure allows us to recover impulse response functions. We find that our estimated impulse responses are consistent with macroeconomic theory and prior studies that focused on one shock at a time. We then use those identified shocks to inquire into the origins of business cycles in the U.S. Identified shocks that originate in household behavior, government expenditure, or monetary policy decisions ("demand" shocks) account for a much larger proportion of business cycles then shocks associated with technical progress or energy costs ("supply shocks"). U.S. business cycles thus have their origins more in demand fluctuations rather than supply shocks.

References

Amir-Ahmadi, P. and T. Drautzburg (2017, May). Identification Through Heterogeneity. Working Papers 17-11, Federal Reserve Bank of Philadelphia.

Angeletos, G.-M., F. Collard, and H. Dellas (2018, July). Business Cycle Anatomy. Nber working papers, National Bureau of Economic Research, Inc.

Bai, J. and S. Ng (2008, June). Large Dimensional Factor Analysis. *Foundations and Trends(R) in Econometrics 3*(2), 89–163.

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

Beraja, M., E. Hurst, and J. Ospina (2016, February). The Aggregate Implications of Regional Business Cycles. NBER Working Papers 21956, National Bureau of Economic Research, Inc.

Bernanke, B. S., J. Boivin, and P. Eliasz (2005). Measuring the effects of monetary policy: a factor-augmented vector autoregressive (favar) approach. *The Quarterly journal of economics 120*(1), 387–422.

Boivin, J., M. P. Giannoni, and I. Mihov (2009). Sticky prices and monetary policy: Evidence from disaggregated us data. *American economic review 99*(1), 350–84.

Caldara, D. and E. Herbst (2016, May). Monetary Policy, Real Activity, and Credit Spreads : Evidence from Bayesian Proxy SVARs. Finance and Economics Discussion Series 2016-049, Board of Governors of the Federal Reserve System (US).

Canova, F. and G. D. Nicolo (2002, September). Monetary disturbances matter for business fluctuations in the G-7. *Journal of Monetary Economics 49*(6), 1131–1159.

Carter, C. K. and R. Kohn (1994, 09). On Gibbs sampling for state space models. *Biometrika 81*(3), 541–553.

Chudik, A. and M. H. Pesaran (2016, February). Theory And Practice Of Gvar Modelling. *Journal of Economic Surveys 30*(1), 165–197.

De Graeve, F. and A. Karas (2010, July). Identifying VARs through Heterogeneity: An Application to Bank Runs. Working paper series, Sveriges Riksbank (Central Bank of Sweden).

Durbin, J. and S. J. Koopman (2012). *Time Series Analysis by State Space Methods.* Oxford University Press.

Faust, J. (1998). The robustness of identified VAR conclusions about money. *Carnegie-Rochester Conference Series in Public Policy 49*, 207–244.

Fernald, J. (2014). A quarterly, utilization-adjusted series on total factor productivity. Federal Reserve Bank of San Francisco.

Fulford, S. L. and F. Schwartzman (2015). The benefits of commitment to a currency peg: Aggregate lessons from the regional effects of the 1896 u.s. presidential election. Technical report, Working Paper 13-18, Federal Reserve Bank of Richmond.

Gelman, A., J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin (2013). *Bayesian Data Analysis*. Chapman & Hall.

Geweke, J. (2005). *Contemporary Bayesian Econometrics and Statistics*. Wiley.

Heckler, D. E. (2005). High-technology employment: a naics-based update. *Monthly Lab. Rev. 128*, 57.

Holly, S. and I. Petrella (2012, November). Factor Demand Linkages, Technology Shocks, and the Business Cycle. *The Review of Economics and Statistics 94*(4), 948–963.

Justiniano, A., G. E. Primiceri, and A. Tambalotti (2010). Investment shocks and business cycles. *Journal of Monetary Economics 57*(2), 132–145.

Kilian, L. (2009, June). Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market. *American Economic Review 99*(3), 1053–1069.

Kociecki, A. (2010). A Prior for Impulse Responses in Bayesian Structural VAR Models. *Journal of Business & Economic Statistics 28*(1), 115–127.

Koop, G. and D. Korobilis (2010, July). Bayesian Multivariate Time Series Methods for Empirical Macroeconomics. *Foundations and Trends(R) in Econometrics 3*(4), 267–358.

Mertens, K. and M. O. Ravn (2013, June). The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States. *American Economic Review 103*(4), 1212–1247.

Nakamura, E. and J. Steinsson (2008). Five facts about prices: A reevaluation of menu cost models. *The Quarterly Journal of Economics 123*(4), 1415–1464.

Pasten, E., R. Schoenle, and M. Weber (2018). The propagation of monetary policy shocks in a heterogeneous production economy. Technical report, National Bureau of Economic Research.

Rajan, R. G. and L. Zingales (1998). Financial dependence and growth. *The American Economic Review 88*(3), 559–586.

Ramey, V. A. (2011). Identifying Government Spending Shocks: It's all in the Timing. *The Quarterly Journal of Economics 126*(1), 1–50.

Romer, C. D. and D. H. Romer (2004, September). A New Measure of Monetary Shocks: Derivation and Implications. *American Economic Review 94*(4), 1055–1084.

Rubin, D. B. (1984, 12). Bayesianly justifiable and relevant frequency calculations for the applied statistician. *The Annals of Statistics 12*(4), 1151–1172.

Rubio-Ramirez, J. F., D. F. Waggoner, and T. Zha (2010). Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference. *Review of Economic Studies 77*(2), 665–696.

Schwartzman, F. (2014). Time to produce and emerging market crises. *Journal of Monetary Economics 68*, 37–52.

Sims, C. A. and T. Zha (1998, November). Bayesian Methods for Dynamic Multivariate Models. *International Economic Review 39*(4), 949–968.

Stock, J. H. and M. W. Watson (2005). Implications of dynamic factor models for var analysis. Technical report, National Bureau of Economic Research.

Uhlig, H. (2005, March). What are the effects of monetary policy on output? Results from an agnostic identification procedure. *Journal of Monetary Economics 52*(2), 381–419.

Zeev, N. B. and E. Pappa (2017, August). Chronicle of a War Foretold: The Macroeconomic Effects of Anticipated Defence Spending Shocks. *Economic Journal 127*(603), 1568–1597.

Appendices

In the following appendices we show plots of our data, and present various robustness checks as well as additional impulse response plots. To economize on space, we focus on presenting mainly the variance decomposition for business cycle frequencies, which can be directly compared to the table in the paper. Throughout we find that our finding on the importance of demand shocks is robust.

A Data Plots

First, we plot the aggregate data we use in our benchmark analysis.

Figure 10: Aggregate data

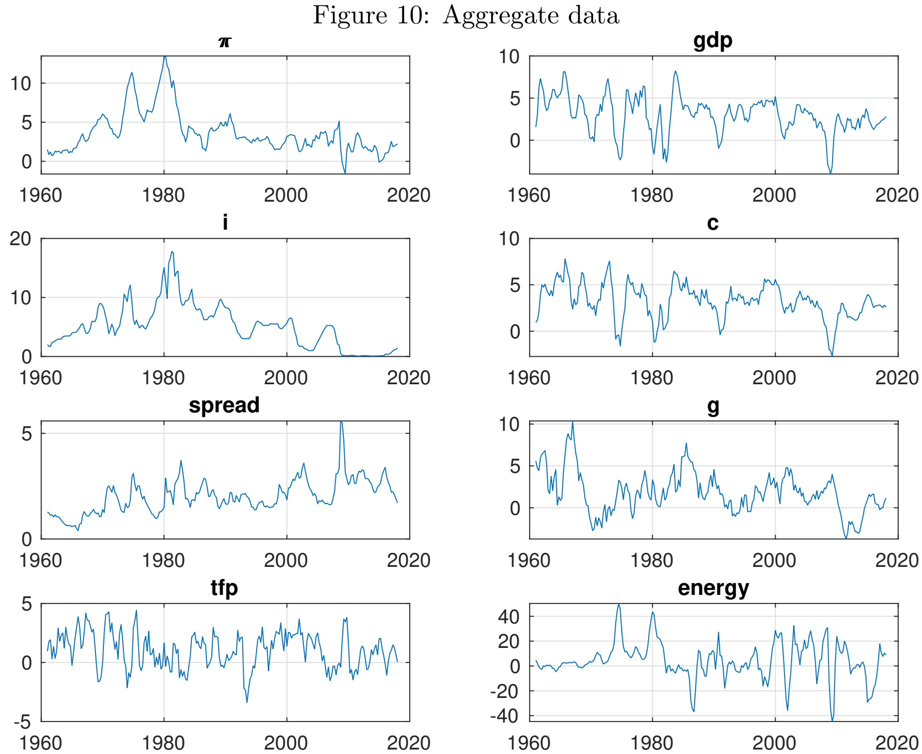

Next, we look at our sectoral data through two different lenses. First, we want to highlight the large dispersion across sectors. To do this, we plot the 5th, 50th, and 9th percentile (across sectors) of our sectoral data over time.

Figure 11: Percentiles of sectoral data

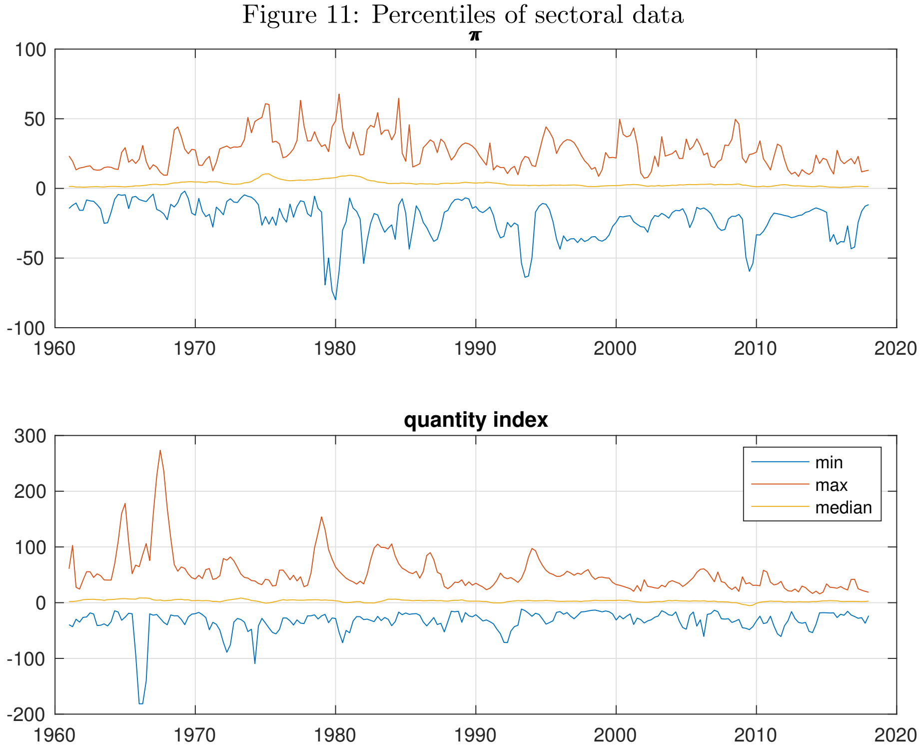

While the previous figure conveys the large dispersion across sectors in outcomes over time, it does not directly speak to how much information is used to identify structural shocks. To do this, we next plot for each of our sectoral indicators that we use to identify shocks, the average time series across all sectors in the top and bottom 10 percent of the distribution according to the different sectoral indicators. In our identification approach, each of the sectors in those parts of the distribution is used separately to identify the shocks. Averaging them provides an informal lower bound on the variation that we exploit for identification.

Figure 12: Sectoral data for identification

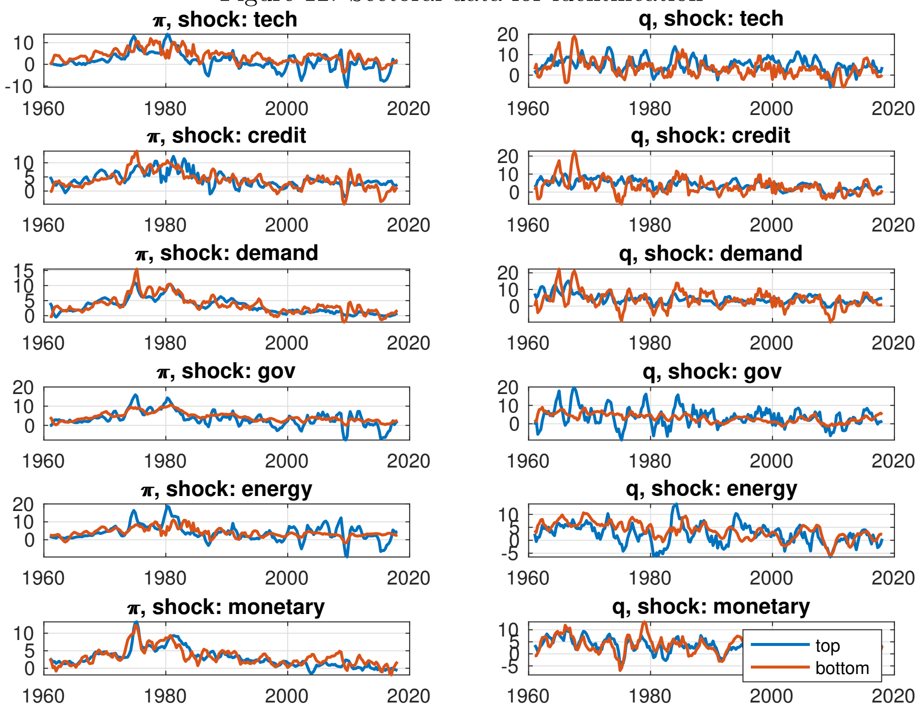

B Using Instruments for Shocks

Recently, there has been substantial interest in using external information/instruments for structural shocks to help identification of the effects of these shocks (Mertens and Ravn (2013)). Here we borrow ideas from Caldara and Herbst (2016) to incorporate these instruments in our Bayesian framework. to do so, we estimate equations of the following form:

$$m_t^i = m^i + a^i \varepsilon_t^i + u_t^i$$

where m_t^i is an instrument for the ith elements of ε_t. m^i captures any possible differences in means across the instrument and true shock, whereas a^i and the variance of u_t^i (which we assume to be a Gaussian iid shock) capture how informative the instrument is by determining the signal to noise ratio in the instrument. Once the parameters are estimated in a separate Gibbs sampling step, we add the instrument equations above to the state space system that is used to generate draws of ε_t. We use four instruments:

1. the government spending news shock from Ramey (2011)
2. the government spending news shock from Zeev and Pappa (2017)
3. the monetary shock from Romer and Romer (2004)
4. the exogenous oil price shock from Kilian (2009)

	tech	credit	demand	gov	energy	monetary
π	6.9	10.6	23.9	7.1	10.7	34.2
gdp	6.4	15.4	25.8	9.1	6.5	28.3
i	6.2	9.2	31.7	4.8	10.7	29.6
c	2.6	7.1	46.1	2.9	3.2	33.9
spread	4.2	30.0	13.5	7.7	8.1	29.7
g	5.9	11.8	18.2	23.1	6.1	28.1
tfp	13.4	13.2	13.2	23.9	5.9	23.7
energy	5.9	10.7	21.1	9.7	11.9	33.2

Table 4: Mean of variance decomposition across business cycle frequencies and posterior draws, using instruments

We truncate our sample to the largest time period so that all shocks are available. Monthly series are averaged to quarterly values. As can be seen from table 4 our results actually become slightly stronger. Unfortunately, the instruments themselves do not add substantial information, as the posteriors for a^i center around 0 and the estimated standard deviation for u_t^i is large for all instruments.

C Alternative Timing in Sectors

In our benchmark we assume that

$$X_t^i = \mu^i + \sum_{l=1}^{L^X} B_l^i X_{t-l}^i + \sum_{l=1}^{L^Y} C_h^i Y_{t-l} + D^i \varepsilon_t + w_t^i \tag{10}$$

We could instead assume that

$$X_t^i = \mu^i + \sum_{l=1}^{L^X} B_l^i X_{t-l}^i + \sum_{l=1}^{L^Y} C_h^i Y_{t-l+1} + D^i \varepsilon_t + w_t^i \tag{11}$$

so that the contemporaneous aggregate variable enters in the sectoral equations.

	tech	credit	demand	gov	energy	monetary
π	5.5	18.6	12.3	6.8	21.2	22.5
gdp	1.8	35.8	2.3	1.6	0.5	57.1
i	9.0	23.2	10.8	5.4	3.6	39.3
c	3.3	12.9	64.4	2.7	1.7	11.4
spread	2.1	45.9	9.0	2.9	2.1	35.7
g	12.7	33.6	16.4	9.1	2.4	24.4
tfp	44.2	31.4	5.1	0.9	14.4	3.2
energy	0.5	2.1	0.8	0.4	93.9	1.4

Table 5: Mean of variance decomposition across business cycle frequencies and posterior draws, alternative timing in sectors

D Using more sectors to identify shocks

In our benchmark analysis, we use the top and bottom 10 percent of sectors (as ranked by the various indicators). Here we use the top and bottom 25 percent instead.

	tech	credit	demand	gov	energy	monetary
π	10.2	18.7	16.1	5.1	3.5	37.7
gdp	13.0	27.9	19.4	4.1	5.7	20.8
i	8.3	11.3	27.4	3.9	3.9	36.1
c	14.4	21.6	29.6	2.9	6.3	15.6
spread	10.6	37.2	11.4	2.9	4.4	25.7
g	11.0	26.9	14.9	12.6	7.5	19.2
tfp	25.2	18.1	15.0	3.1	4.3	17.0
energy	9.0	18.7	14.4	4.9	14.1	28.6

Table 6: Mean of variance decomposition across business cycle frequencies and posterior draws, more sectors used in identification

E A Labor Market Shock

In this section we add one additional element to ε_t relative to our benchmark: a labor market shock, which we identify as a shock that has a larger impact on sectors where compensation to employees as a share of value added is higher.

	tech	credit	demand	gov	energy	monetary	labor
π	6.8	7.3	7.5	5.7	8.4	50.3	5.7
gdp	9.6	13.8	20.3	5.9	10.0	21.2	8.8
i	7.1	8.3	15.8	5.6	12.0	35.0	6.8
c	15.7	11.0	27.9	6.1	8.8	17.5	6.9
spread	11.8	28.0	8.4	4.6	8.3	21.4	11.6
g	6.3	12.9	15.6	15.3	10.9	20.7	9.4
tfp	32.5	7.5	9.6	4.0	20.8	11.3	4.9
energy	7.3	8.9	8.7	5.4	26.6	29.1	6.1

Table 7: Mean of variance decomposition across business cycle frequencies and posterior draws, with labor shock

F Lag Length

In our benchmark analysis we assumed $L = L^X = 6$ lags. Here instead we assume $L = L^X = 4$.

	tech	credit	demand	gov	energy	monetary
π	9.0	14.2	7.4	14.3	10.9	27.6
gdp	12.9	21.1	6.0	9.0	20.1	16.6
i	11.5	16.0	8.2	9.0	14.5	25.5
c	21.4	21.0	10.2	6.7	16.2	15.4
spread	18.0	40.0	3.7	5.4	11.3	15.7
g	18.3	27.3	4.2	15.1	13.1	10.0
tfp	32.8	23.9	2.8	3.0	9.2	21.3
energy	10.7	17.6	5.3	9.0	25.7	18.0

Table 8: Mean of variance decomposition across business cycle frequencies and posterior draws, 4 lags

G More Elements in ε without prior restrictions

In our benchmark model, we add 3 elements to ε_t for which we do not impose any prior information. To check whether the choice of 3 additional shocks is crucial, we now present in teble 9 results with 10 additional shocks.

	tech	credit	demand	gov	energy	monetary
π	11.5	11.8	14.4	5.8	14.3	17.9
gdp	10.9	10.0	19.9	3.4	16.3	10.1
i	9.3	8.9	19.2	3.6	17.8	12.5
c	10.6	11.4	26.5	3.5	15.2	12.9
spread	6.4	36.4	12.2	3.6	16.9	8.9
g	10.7	10.2	24.3	7.6	12.0	7.8
tfp	43.4	7.0	14.5	2.1	13.1	7.0
energy	15.4	13.5	13.3	4.2	23.5	12.5

Table 9: Mean of variance decomposition across business cycle frequencies and posterior draws, 10 additional shocks

H Additional Impulse Responses

H.1 Plots with 5th and 95th Percentile Bands

Figure 13: Responses to Technology Shock

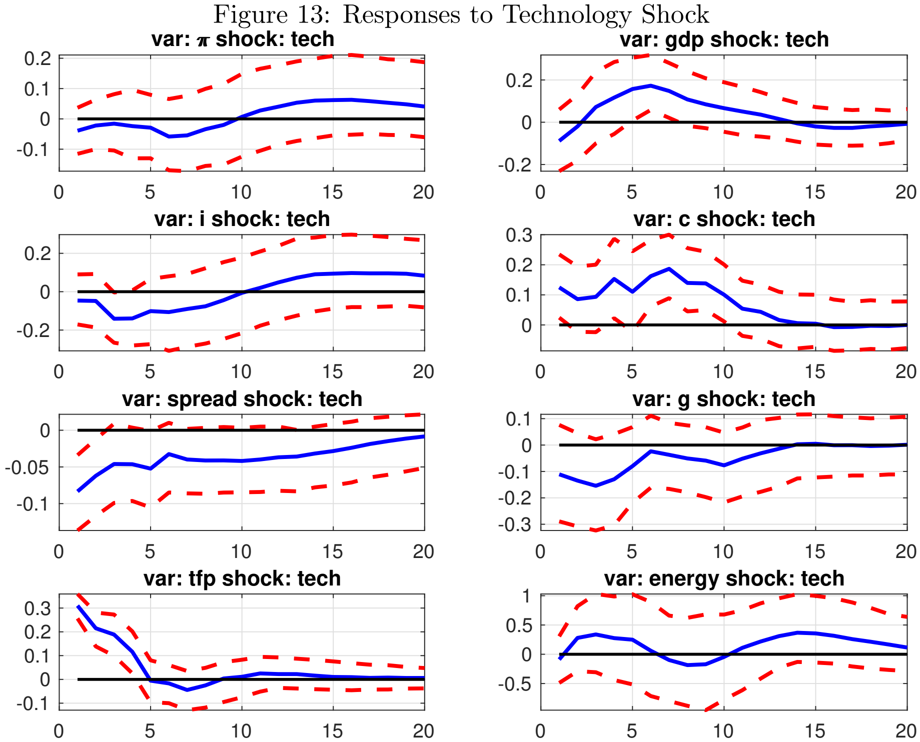

Figure 14: Responses to Corporate credit Shock

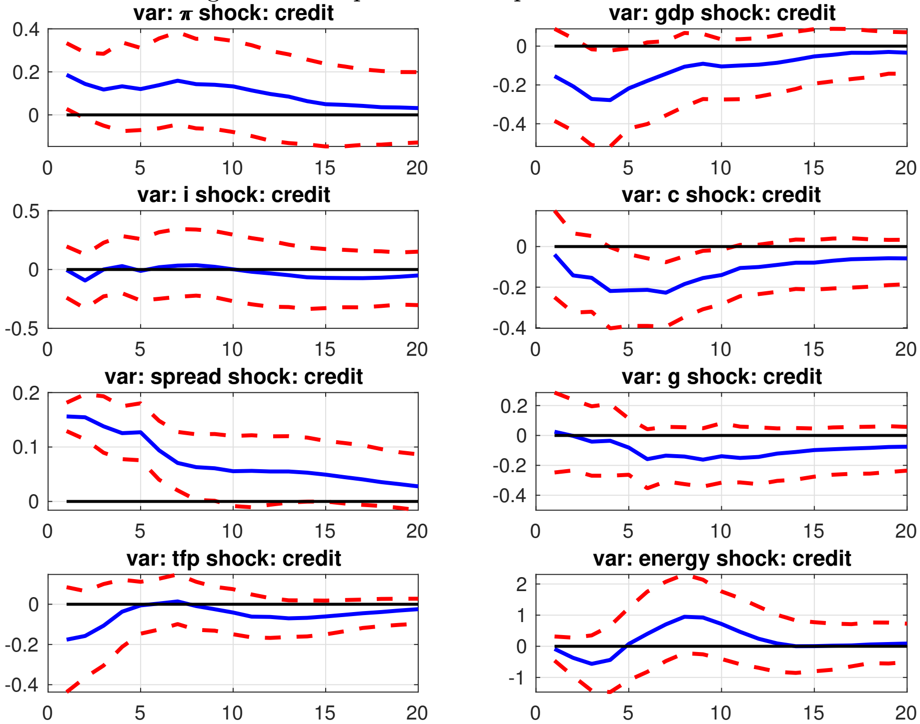

Figure 15: Responses to Government Shock

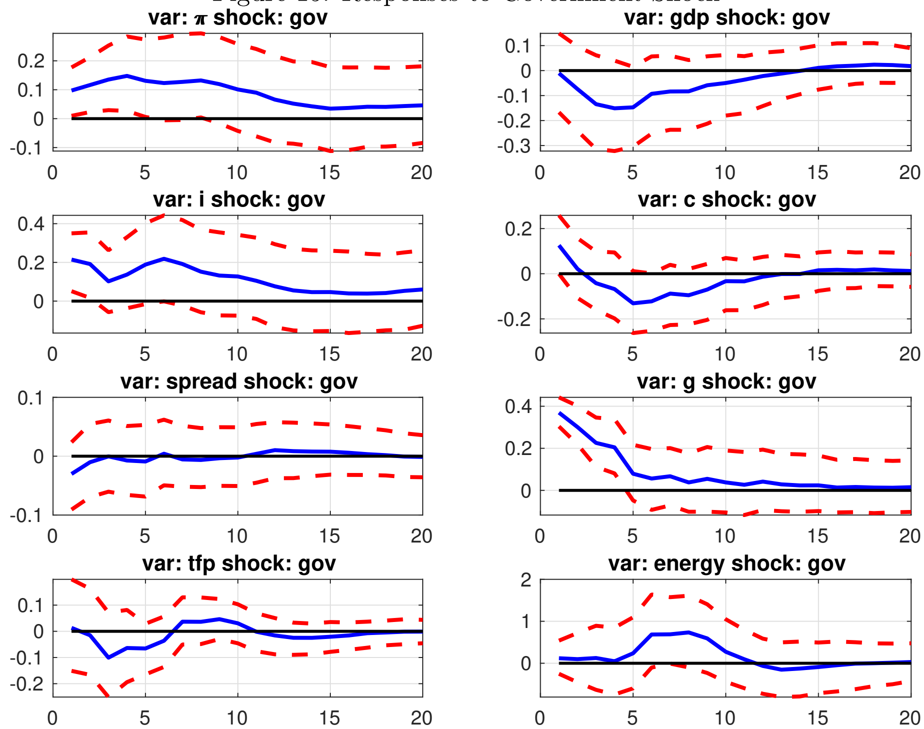

Figure 16: Responses to Energy Shock

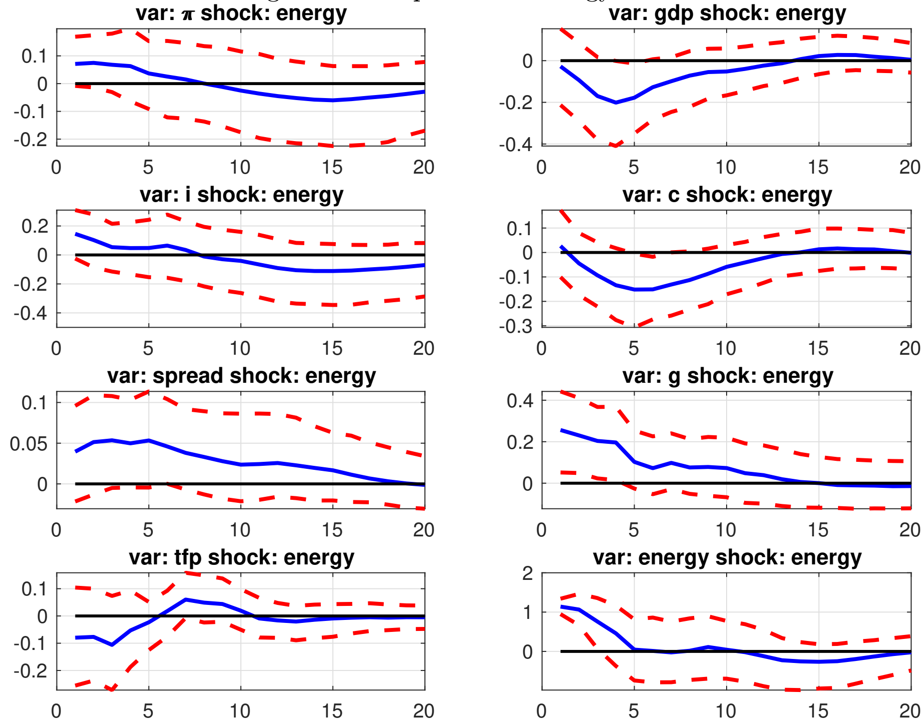

H.2 Digging Deeper into IDentification Assumptions for other Variables

Figure 17: Responses to Technology Shocks: Comparison of Identification Schemes

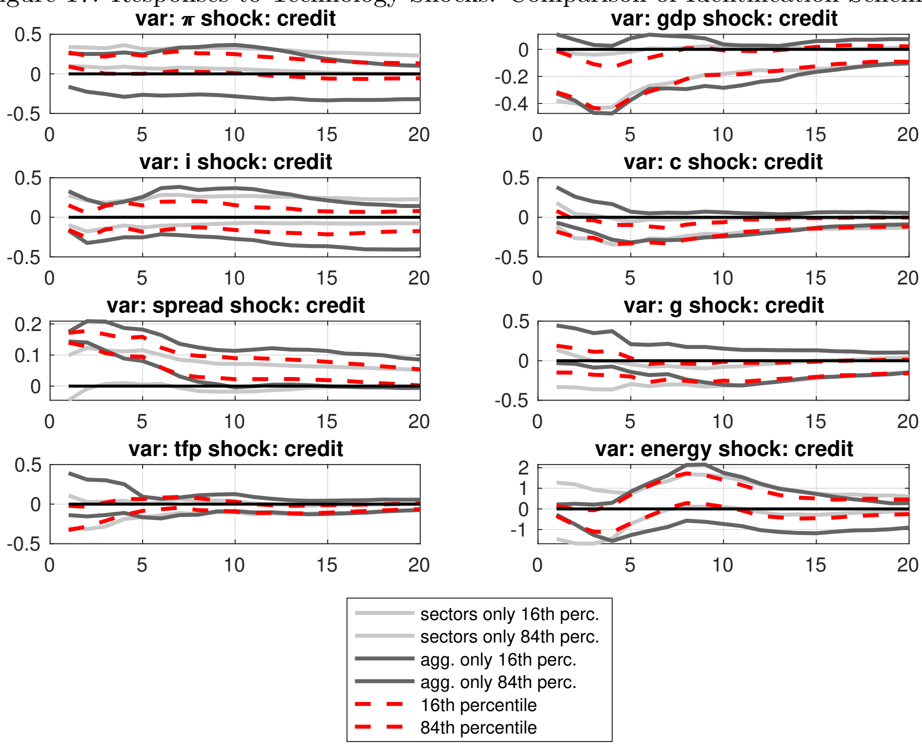

Figure 18: Responses to Credit Shocks: Comparison of Identification Schemes

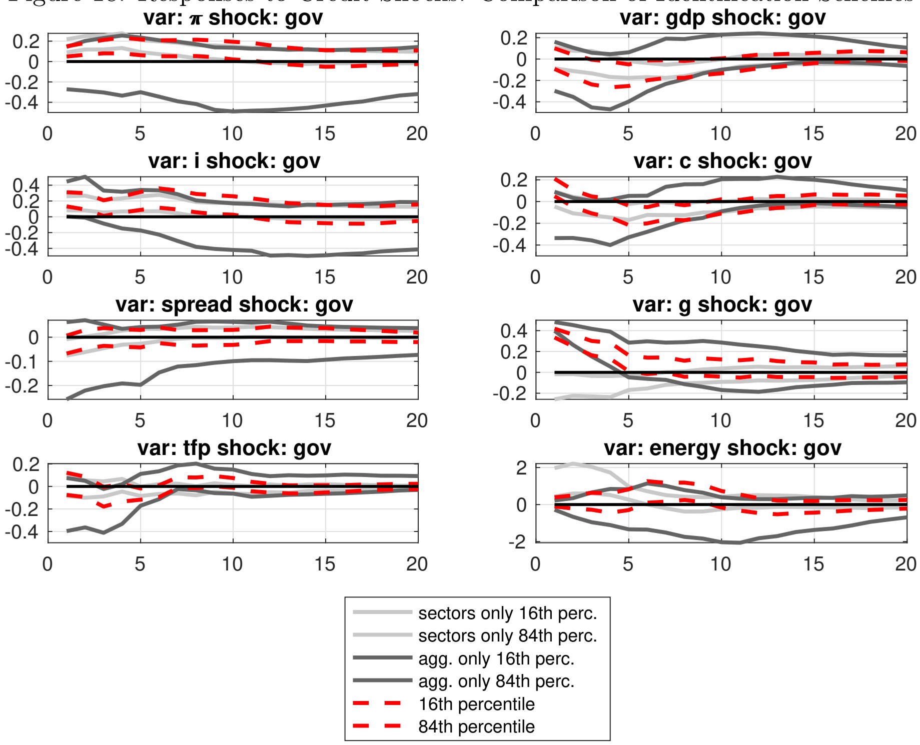

Figure 19: Responses to Government Shocks: Comparison of Identification Schemes

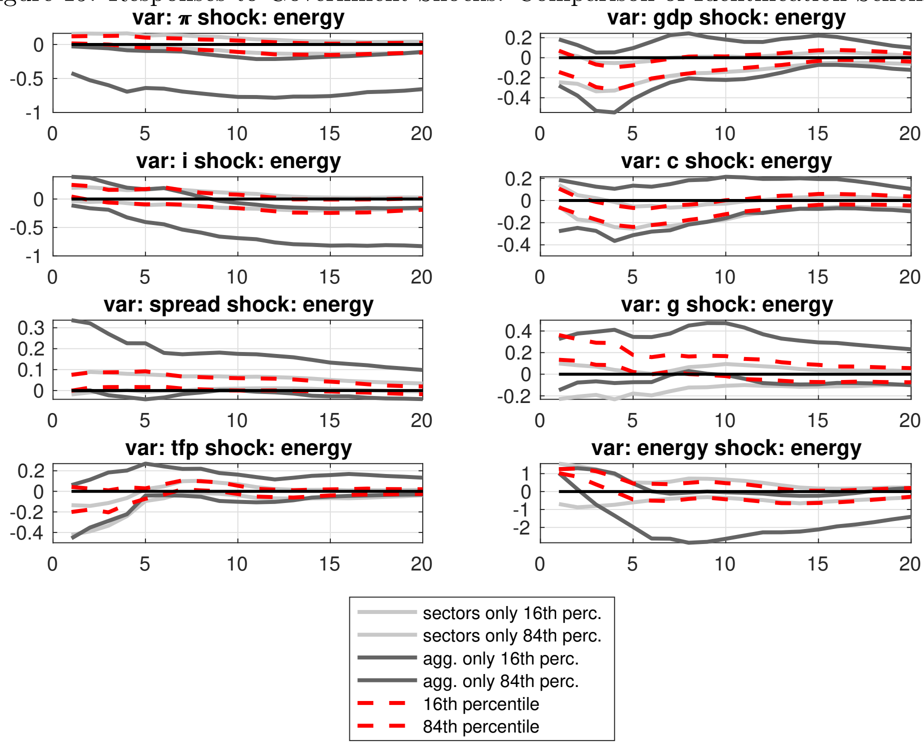

Figure 20: Responses to Energy Shocks: Comparison of Identification Schemes

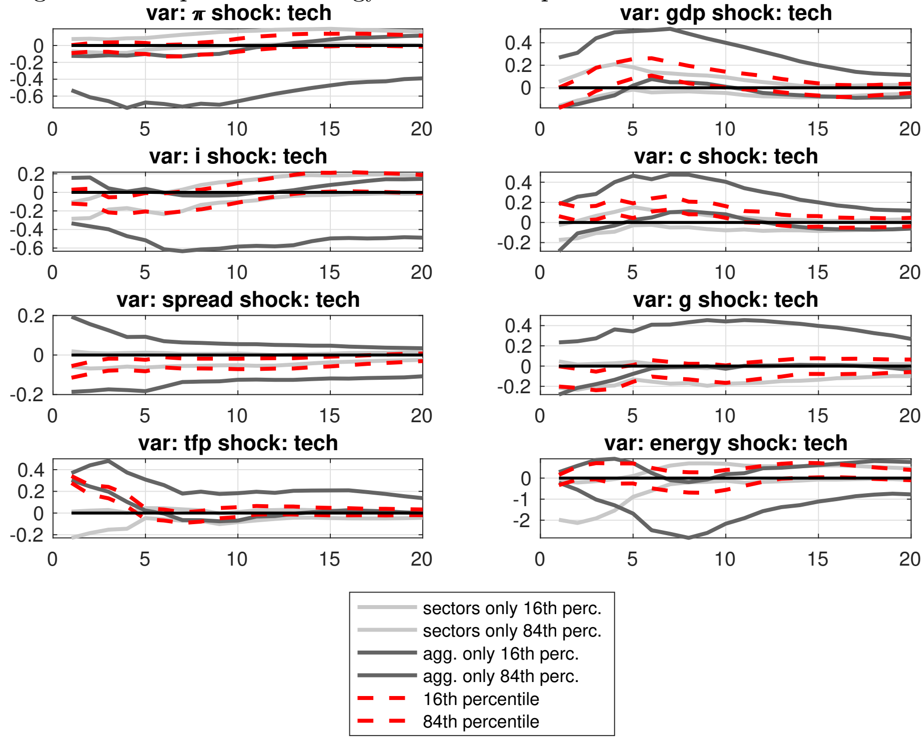

I Displaying Error Bands for Our Monte Carlo Exercise

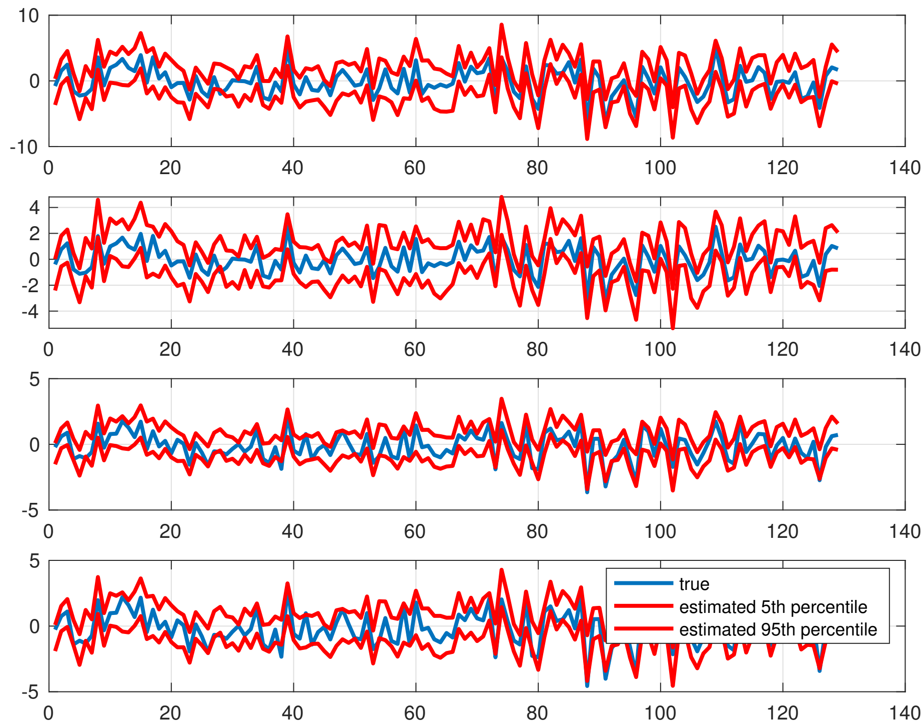

Figure 21: True effect and 90 percent estimated error bands