Letting Different Views about Business Cycles Compete

Web-Appendix

Paul Beaudry, University of British Columbia
Bernd Lucke, University of Hamburg

The goal of this appendix is to explore whether the identification strategies used in our paper "Letting Different Views about Business Cycles Compete" can plausibly uncover structural impulse responses. To this end, we build a model which theoretically satisfies the type of identifying restrictions discussed in the paper, and then we use data generate from the model to see whether our identification strategy would actually recover the structural stocks. In order to simplify this exercise, we build a model driven by only four real shocks: innovations in TFP, innovations in investment specific technological change, news about future innovations in TFP and preferences shocks. We omit from the model a monetary shock and nominal rigidities.

1. The Model

The representative agent maximizes

$$U = E_0 \sum_{t=0}^{\infty} \beta^t \left(\ln C_t + \ln\left(1 - \psi_t N_t\right)\right)$$

where $0 < \beta < 1$, C_t is consumption, N_t is hours and ψ_t is a stationary preference shock with data generating process $\ln \psi_t = \rho_\psi \ln \psi_{t-1} + \varepsilon_{\psi,t}$.

The resource constraint is

$$A_t \left(u_t K_{t-1}\right)^{1-\alpha} N_t^{\alpha} = C_t + \frac{I_t}{z_t}. \tag{1}$$

Here u_t is the capacity utilization rate, K_t and I_t are capital and investment, respectively, measured in constant base year quality, A_t is total factor productivity (TFP), z_t is investment specific technology (IST) and $0 < \alpha < 1$.

The depreciation rate δ is an increasing and convex function of the capital utilization rate, $\delta(u) = \delta_0 u^{\omega}, \quad \omega > 1$. Let γ_K be the deterministic steady state growth factor of capital and δ the steady state depreciation rate. A function $\phi_K, \quad \phi_K(1) = \phi_K{}'(1) = 0, \quad \phi_K{}''(1) > 0$ describes nonzero capital adjustment costs outside of the steady state:

$$K_t = \left(1 - \delta(u_t)\right) K_{t-1} + I_t \left(1 - \phi_K\left(\frac{I_t}{\left(\gamma_K - 1 + \delta\right) K_{t-1}}\right)\right) \tag{2}$$

Log IST is a random walk with drift. Log TFP is a unit root process subject to a surprise technology innovation $\varepsilon_{A,t}$ and an anticipated innovation (news shock) $\varepsilon_{N,t-d}$.

$$\ln A_t = \ln \gamma_A + \ln A_{t-1} + \varepsilon_{A,t} + \varepsilon_{N,t-d}$$
$$\ln z_t = \ln \gamma_z + \ln z_{t-1} + \varepsilon_{z,t}$$

The news shock becomes publicly known *d* periods before the innovation actually affects TFP. For simplicity, we assume that the effect of $\varepsilon_{N,t-d}$ on TFP materializes fully within a single period, *t*. A more realistic treatment would have this effect diffuse over several periods.

$\varepsilon_{A,t}$, $\varepsilon_{z,t}$ and $\varepsilon_{N,t}$ are mutually independent white noise processes with variances σ_A^2, σ_z^2, σ_N^2, respectively.

We derive the first order conditions and compute a measure of stock prices as the value of installed capital in consumption units, i. e. the shadow price of capital divided by the shadow price of consumption.

2. Solution and Calibration:

We stationarize the system by dividing each variable by its stochastic growth trend, i. e. the appropriate combinations of

$$A_t = \gamma_A A_{t-1} \exp\left(\varepsilon_{A,t} + \varepsilon_{N,t-d}\right)$$
$$z_t = \gamma_z z_{t-1} \exp\left(\varepsilon_{z,t}\right)$$

We use Uhlig's (1997) method of undetermined coefficients to find the equilibrium state space representation of the model. The news shock is incorporated in the law of motion for the driving processes by specifying

$$\begin{pmatrix} \varepsilon_{z,t} \\ \hat{\psi}_t \\ \varepsilon_{TFP,t} \\ \varepsilon_{N,t} \\ \varepsilon_{N,t-1} \\ \vdots \\ \varepsilon_{N,t-d+1} \end{pmatrix} = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & \rho_\psi & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & \cdots & 1 \\ 0 & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 1 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} \varepsilon_{z,t-1} \\ \hat{\psi}_{t-1} \\ \varepsilon_{TFP,t-1} \\ \varepsilon_{N,t-1} \\ \varepsilon_{N,t-2} \\ \vdots \\ \varepsilon_{N,t-d} \end{pmatrix} + \begin{pmatrix} \varepsilon_{z,t} \\ \varepsilon_{\psi,t} \\ \varepsilon_{A,t} \\ \varepsilon_{N,t} \\ 0 \\ \vdots \\ 0 \end{pmatrix}$$

We calibrate the parameters of the model as follows:

$$\alpha = 0.64, \quad \beta = 0.985, \quad \delta = 0.025, \quad \delta_0 = 0.035, \quad \omega = 1.15, \quad \gamma_A = \gamma_z = 1.002, \quad \rho_\psi = 0.5, \quad \phi_K'' = 1.3$$

We ran simulations for three different values for the delay parameter d, $d \in \{1,4,8\}$. We document only $d = 8$ in this appendix, since this seems to be the most realistic value. Smaller values of d generally make identification easier. We use the variances σ_A^2, σ_z^2, σ_N^2 to calibrate the model such that the variance decomposition of hours is similar to the variance decomposition we find in the real-world data.

3. Simulations:

We set the presample values of the variables expressed as stationary percentage deviations from the stochastic growth path equal to zero. Then we simulate the model over 500 periods. We throw the first 300 observations away to get rid of the influence of the presample values. We compute the simulated stochastic growth paths from the structural residuals and add the appropriate stochastic trend to the endogenous stationary variables.

For each set of calibrated parameters 100 sets of artificial time series with 200 observations are generated. The samples are analysed using the GAUSS sourcecode of JMulti 4.23, cf. www.jmulti.de. We estimate four-dimensional VECMs consisting of the logs of TFP, IST, stock prices and hours, imposing two cointegrating vectors and a lag length in the VAR of d+2. The structural decompositions are computed with identification schemes analogous to ID1, ID2 and ID3 in our paper: we simply delete the fifth row and fifth column in the impact matrices B and L, since the simulated model does not have nominal variables.

For the identified structural residuals we compute the correlation with their true (simulated) counterpart. Below we report the mean absolute correlation for each shock and the respective sample standard error. We confine the simulations to 100 runs per calibrated set of parameters, since some experiments showed that extending this to 1000 runs yields very similar results. We also compute mean and standard deviations over 100 runs of the impulse response functions and of the forecast error variance decompositions (FEVD).

4. Results

We focus on the case d=8. Given our results, this delay of two years for a news shock to affect TFP seems to be appropriate. Smaller values of d generally increase the quality of shock identification.

Structural identification does not allow direct inference on the variances of the structural shocks. Since the performance of the identification scheme hinges critically on the relative standard deviations of the structural shocks, we experimented with different sets of variances to find a configuration such that our model implies a variance decomposition of hours across the business cycle frequencies which resembles the one we find in the data. We illustrate our procedure in Figures 1-3, whose parameterizations correspond to the following setting:

	Relative standard deviations of			
	TFP shock (pink)	IST shock (red)	News shock (blue)	Preference shock (green)
Figure 1	1	1	1	1
Figure 2	1	1	1	0.02
Figure 3	0.1	0.1	1	0.02

We find that equal standard deviations, Figure 1, attribute much more hours variance to the preference shock than we find in the data. Reducing the relative standard deviations of preference shocks to 2%, cf. Figure 2, results in excessive variance shares of TFP and IST shocks. So we reduce their relative standard deviations to 10% and obtain a variance decomposition of hours similar to our empirical results, cf. Figure 3.

Mean estimated FEVDs		
Figure 1	Figure 2	Figure 3
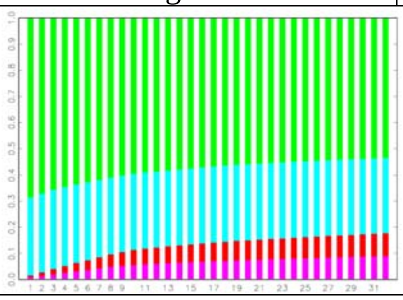	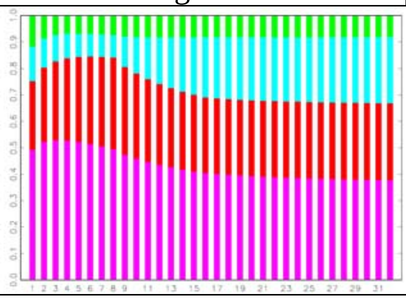	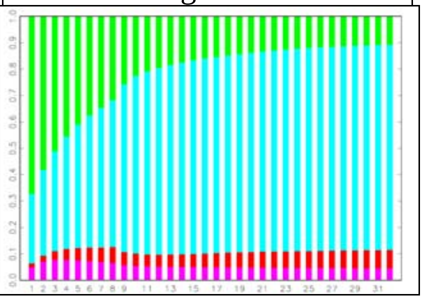

The last setting (recall *d*=8) implies the following mean correlations between identified and true shocks:

Identification ID1 :

		TFP	IST	News	preference shock
Mean abs. correlation	:	0.5604	0.8844	0.8195	0.8095
Standard Deviation	:	0.0457	0.0257	0.1027	0.1094

Identification ID2 :

		TFP	IST	News	preference shock
Mean abs. correlation	:	0.5616	0.7710	0.7958	0.7425
Standard Deviation	:	0.0450	0.1553	0.1091	0.1648

Identification ID3 :

		TFP	IST	News	preference shock
Mean abs. correlation	:	0.5491	0.7085	0.6877	0.7466
Standard Deviation	:	0.0429	0.1877	0.1905	0.1798

Our two main identification schemes ID1 and ID2 identifiy the news shock relatively well (correlations of about 0.8). The identification of preference and IST shock is also satisfactory (correlations between 0.7 and 0.9). More difficult is the identification of the TFP shock (correlation only somewhat more than 0.5). This is not surprising, since the TFP shock is only a minor part of the innovation of the TFP process (standard deviation relative to news shock is only 10%). In fact, reliable identification of TFP shocks seems to be difficult precisely because the TFP shock is unimportant by construction – and it is probably not too problematic if an unimportant shock is not well identified.

Identification ID3 works a little worse than ID1 and ID2. This is particularly true for the news shock, both in terms of mean and standard error. But the ID3 identification in our paper is not designed to identify news shocks well – it is designed to place as little restrictions as possible on IST shocks. This seems to induce a bias towards IST shocks in the identification, since the mean variance decomposition of hours for ID3 clearly overstates the importance of IST shocks and understates the importance of news shocks relative to the true variance decomposition of the model, cf. Figure 3 and Figure 6. Thus, based on our simulations we would not recommend ID3 for assessing the relative importance of different business cycle shocks.

Mean estimated FEVDs

Figure 4 ID1	Figure 5 ID2	Figure 6 ID3

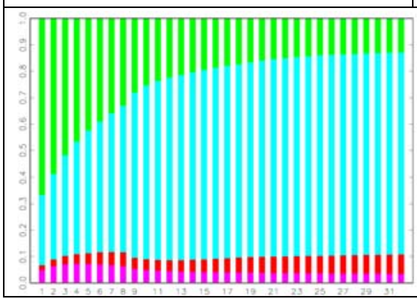
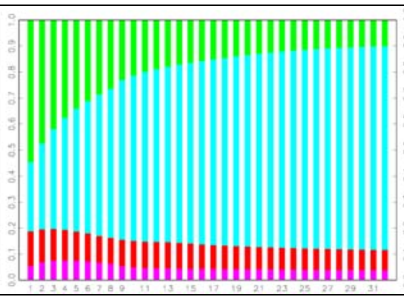
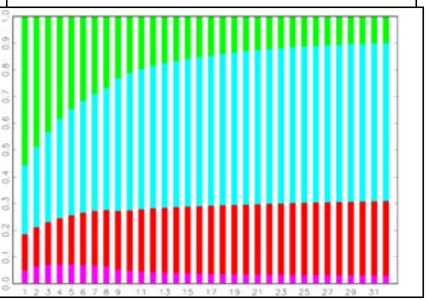

For ID1 and ID2 the mean FEVD seems to be an almost unbiased estimator of the true (model) variance decomposition, cf. Figures 4 and 5. In terms of impulse responses (IR), the mean estimated IR is mostly close to its true model counterpart, but confidence intervals for some IRs can be sizable. Figure 7 shows the mean of the IRs estimated under ID1 (green) along with the model IRs (red) and the estimated ± two standard errors confidence intervals (pink). Results for ID2 and ID3 are very similar. The most relevant aspects to note about this exercise is that our identification strategy retrieves impulse response for the effects on hours of news shocks and preferences shocks which are very close to the theoretical impulse responses implied by the model. While this does not imply that our identification strategies would always work, it does suggests that they have the potential of working quite well.