Summary of Key Papers in the News Literature

Beaudry-Portier AER 2006

(Econometric details presented in class on whiteboard.)

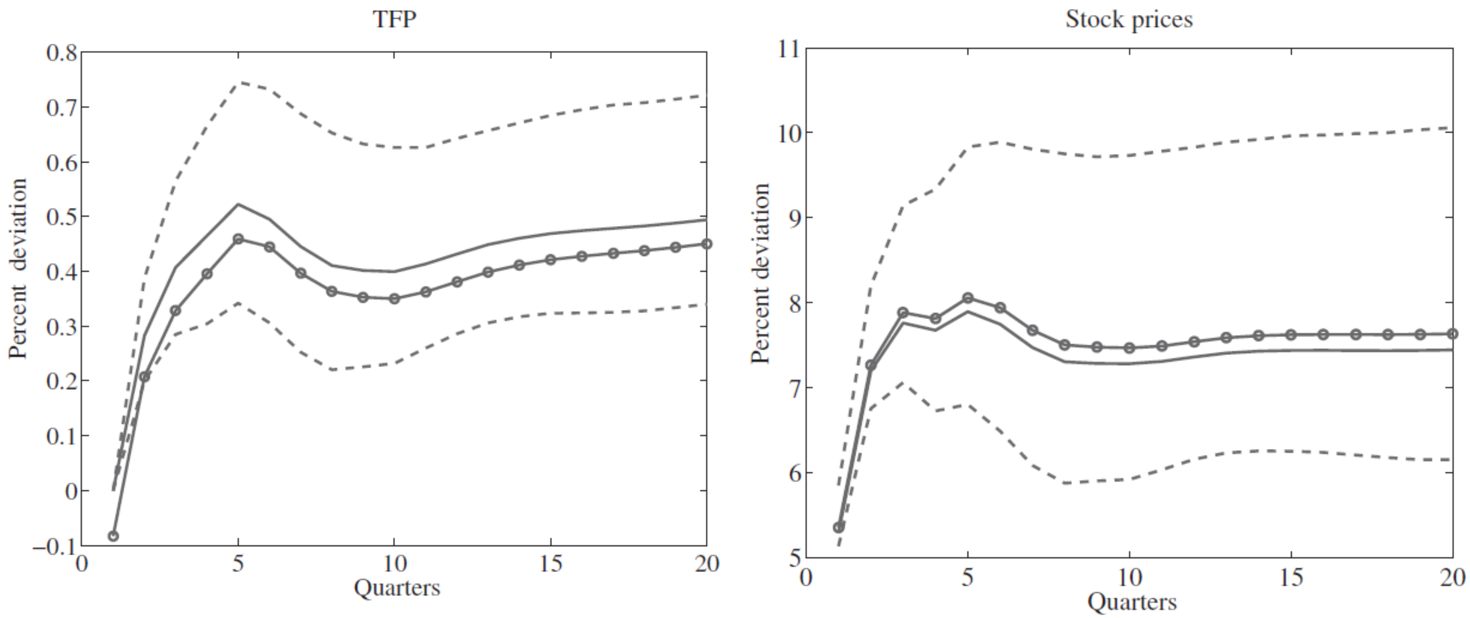

FIGURE 1. IMPULSE RESPONSES TO SHOCKS ε_2 AND $\tilde{\varepsilon}_1$ IN THE (*TFP*, *SP*) VECM

Notes: In both panels of this figure, the bold line represents the point estimate of the responses to a unit ε_2 shock (the shock that does not have instantaneous impact of *TFP* in the short-run identification). The line with circles represents the point estimate of the responses to a unit $\tilde{\varepsilon}_1$ shock (the shock that has a permanent impact on TFP in the long-run identification). Both identifications are done in the baseline bivariate specification (five lags and one cointegrating relation). The unit of the vertical axis is percentage deviation from the situation without shock. Dotted lines represent the 10-percent and 90-percent quantiles of the distribution of the impulse response functions (IRFs) in the case of the short-run identification, this distribution being the Bayesian simulated distribution obtained by Monte-Carlo integration with 2,500 replications, using the approach for just-identified systems discussed in Thomas J. Doan (1992).

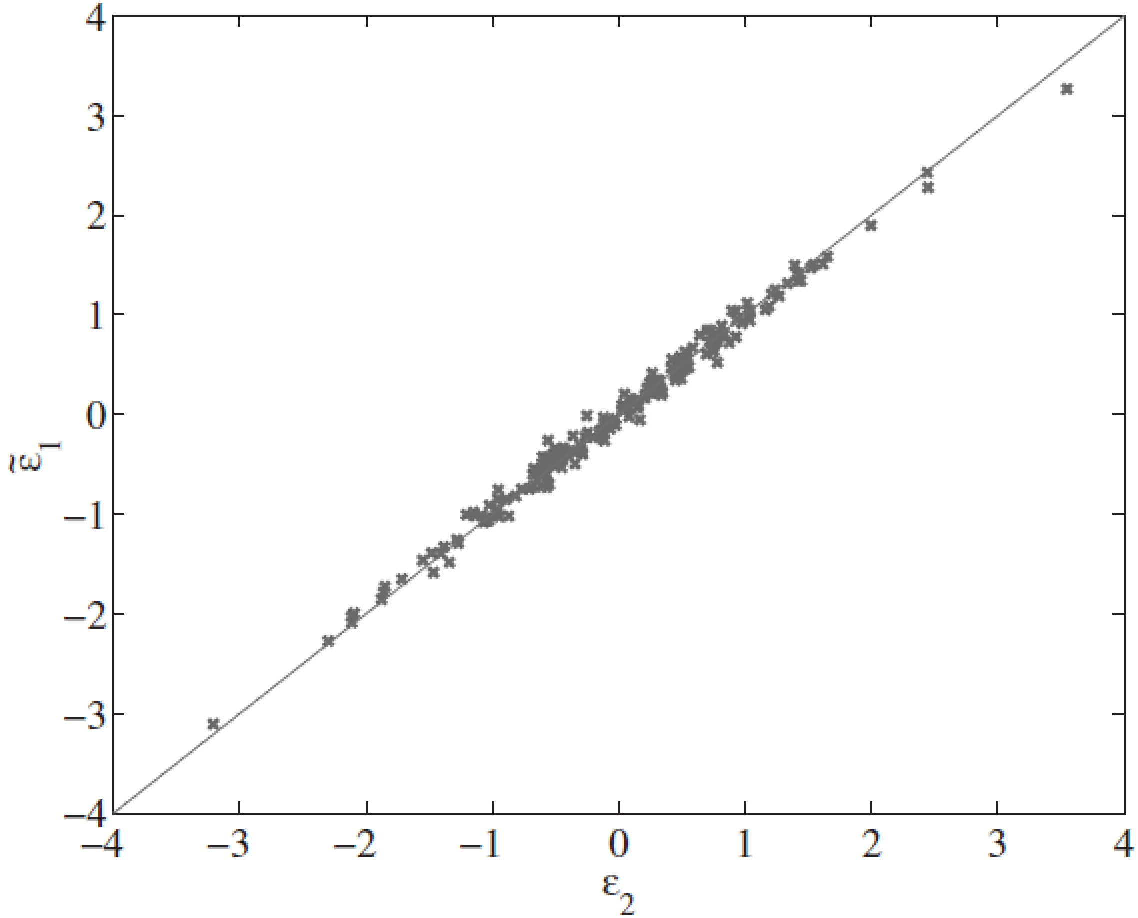

FIGURE 2. PLOT OF ε_2 AGAINST $\tilde{\varepsilon}_1$ IN THE (*TFP*, *SP*) VECM

Notes: This figure plots ε_2 against $\tilde{\varepsilon}_1$. Both shocks are obtained from the baseline bivariate specification (five lags and one cointegrating relation). The straight line is the 45-degree line.

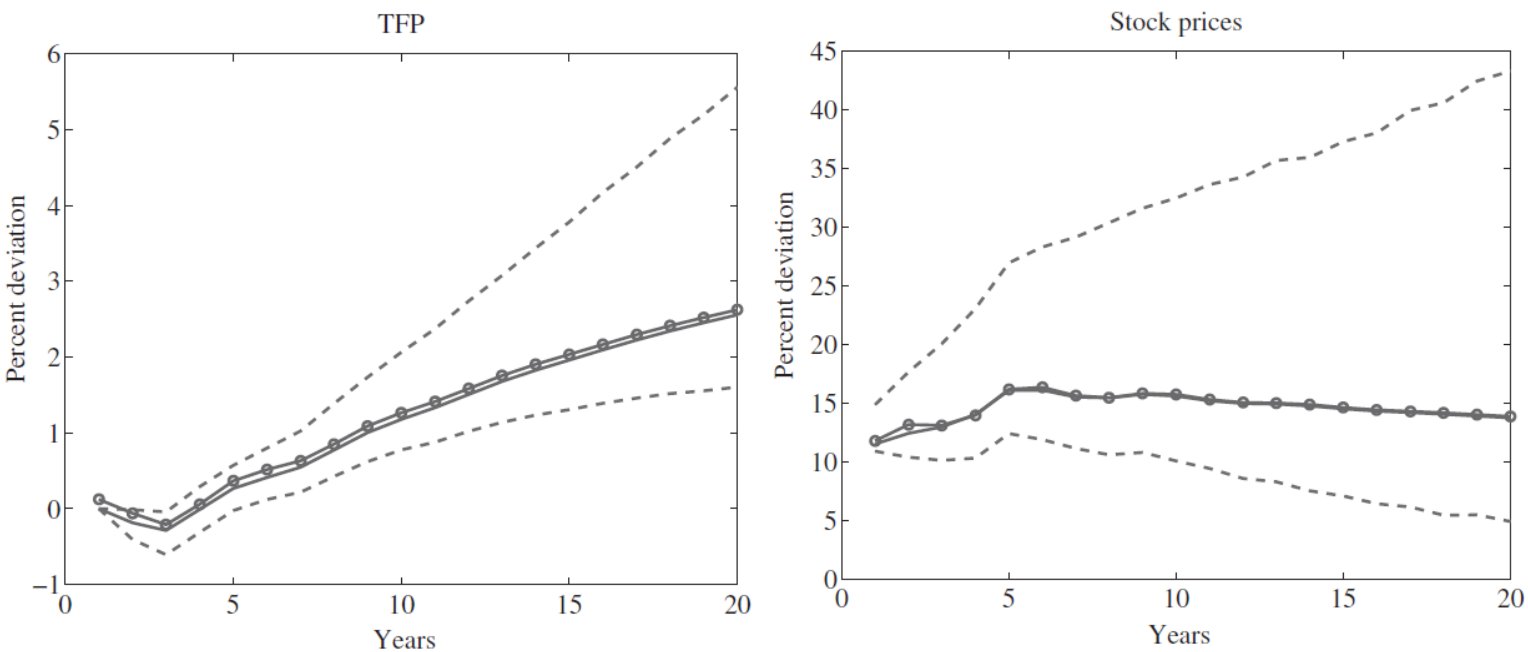

Figure 4. Impulse Responses to Shocks ε_2 and $\tilde{\varepsilon}_1$ in the (*TFP*, *SP*) VECM, Using Basu et al. (2002) Measure of TFP (Annual, 1949–1989)

Notes: In both panels of this figure, the bold line represents the point estimate of the responses to a unit ε_2 shock (the shock that does not have instantaneous impact on *TFP* in the short-run identification). The line with circles represents the point estimate of the responses to a unit $\tilde{\varepsilon}_1$ shock (the shock that has a permanent impact on TFP in the long-run identification). Both identifications are done in the baseline annual specification (two lags and one cointegrating relation). The unit of the vertical axis is percentage deviation from the situation without shock. Dotted lines represent the 10-percent and 90-percent quantiles of the distribution of the *IRF* in the case of the short-run identification, this distribution being the Bayesian simulated distribution obtained by Monte-Carlo integration with 2,500 replications, using the approach for just-identified systems discussed in Doan (1992).

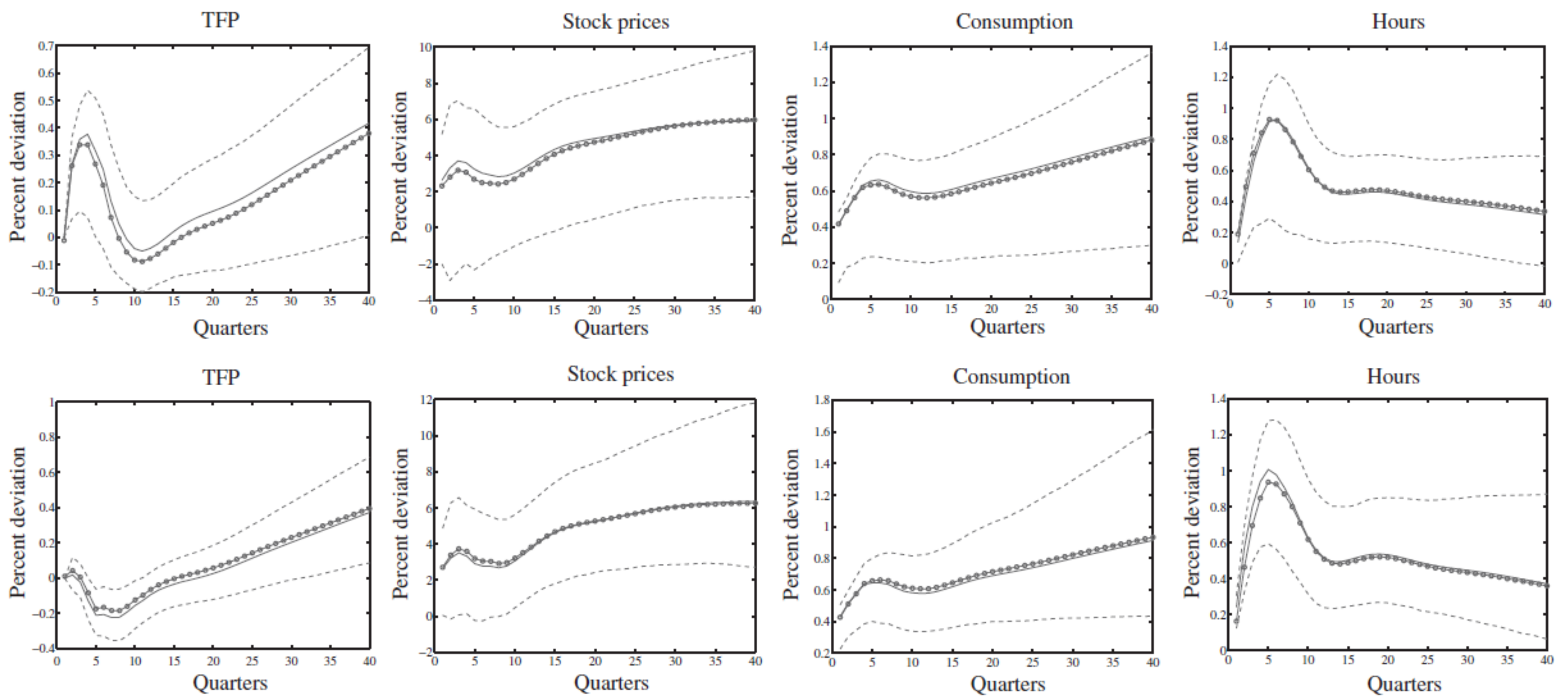

Figure 9. Impulse Responses to ε_2 and $\tilde{\varepsilon}_1$ in the (*TFP*, *SP*, *C*, *H*) VECM, without (upper panels) or with (lower panels) Adjusting TFP for Capacity Utilization

Notes: In each panel of this figure, the bold line represents the point estimate of the responses to a unit ε_2 shock (the shock that does not have instantaneous impact on *TFP* in the short-run identification). The line with circles represents the point estimate of the responses to a unit $\tilde{\varepsilon}_1$ shock (the shock that has a permanent impact on TFP in the long-run identification). In this system with hours, both identifications are done in a specification with five lags and three cointegrating relations, i.e., a VAR in levels. The unit of the vertical axis is percentage deviation from the situation without shock. Dotted lines represent the 10-percent and 90-percent quantiles of the distribution of the IRF in the case of the short-run identification, this distribution being the Bayesian simulated distribution obtained by Monte-Carlo integration with 2,500 replications, using the approach for just-identified systems discussed in Doan (1992).

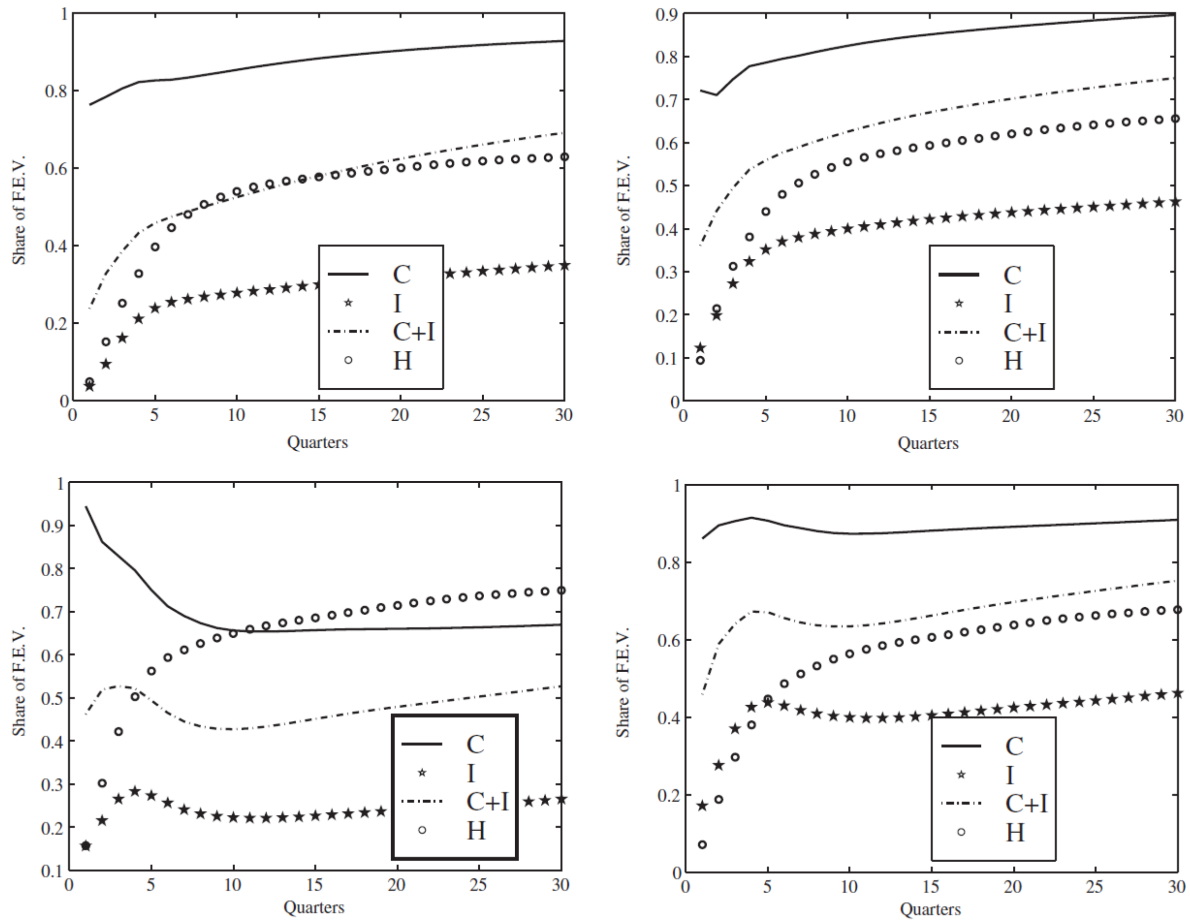

FIGURE 10. SHARE OF THE FORECAST ERROR VARIANCE (F.E.V.) OF CONSUMPTION (C), INVESTMENT I, OUTPUT ($C + I$), AND HOURS (H) ATTRIBUTABLE TO ε_2 (LEFT PANELS) AND TO $\tilde{\varepsilon}_1$ (RIGHT PANELS) IN VECMS, WITH NONADJUSTED TFP (TOP PANELS) OR ADJUSTED TFP (BOTTOM PANELS)

Notes: This figure has four panels. The left panels display the share of the forecast variance of consumption and investment that is attributable to ε_2 (short-run identification) in the (*TFP*, *SP*, *C*, *I*) VECM (five lags and three cointegrating relations), of output ($C + I$) in the (*TFP*, *SP*, *C*, $C + I$) VECM (five lags and three cointegrating relations), and of hours (H) in the (*TFP*, *SP*, *C*, *H*) VECM (five lags and four cointegrating relations, i.e., a VAR in levels). The right panels display the same information in the case of the shock $\tilde{\varepsilon}_1$ (long-run identification). The top row uses a nonadjusted measure of TFP, while TFP is adjusted for variable capacity utilization in the bottom row.

Stock Prices, News, and Economic Fluctuations: Comment

By André Kurmann and Elmar Mertens*

Beaudry and Portier (2006) propose an identification scheme to study the effects of news shocks about future productivity in vector error correction models (VECMs). This comment shows that, when applied to their VECMs with more than two variables, the identification scheme does not have a unique solution. The problem arises from a particular interplay of cointegration assumptions and long-run restrictions.

This comment shows that, in the VECMs with more than two variables estimated by Beaudry and Portier (2006), their identification scheme fails to determine TFP news. Yet these higher-dimension systems are crucial to quantify the business-cycle effects of TFP news.[2] The identification problem arises from the interplay of two assumptions. First, the Beaudry-Portier identification scheme, called BP restrictions from here on, requires that one of the non-news shocks has no permanent impact on either TFP or consumption. Second, the VECMs estimated by Beaudry and Portier (2006) impose that TFP and consumption are cointegrated. This cointegration means that TFP and consumption have the same permanent component, which makes one of the two long-run restrictions redundant and leaves an infinity of candidate solutions. The results reported in Beaudry and Portier (2006) represent just one arbitrary choice among these solutions.[3]

the BP restrictions fail to identify TFP news. The identification scheme and results presented in Beaudry and Portier (2006) therefore do not shed light on the importance of TFP news shocks for business cycles.

John Cochrane, "Shocks," Carnegie-Rochester 1994.

Abstract

What are the shocks that drive economic fluctuations? I examine technology and money shocks in some detail, and briefly review the evidence on oil price and credit shocks. I conclude that none of these popular candidates accounts for the bulk of economic fluctuations. I then examine whether "consumption shocks," reflecting news that agents see but we do not, can account for fluctuations. I find that it may be possible to construct models with this feature, though it is more difficult than is commonly realized. If this view is correct, we will forever remain ignorant of the fundamental causes of economic fluctuations.

We have examined popular candidates for shocks, and found little solid evidence that they account for the bulk of business-cycle fluctuations. Shocks to consumption, output, or other endogenous variables dominate most calculations. Other contenders, such as government spending or financing shocks, are not quantitatively plausible.

Since we can not seem to find observable exogenous shocks, how about *unobservable* shocks? Surely agents have much more information than we do. Suppose they get bad news about the future. Then, consumption declines and sets off a recession. We economists, like Hall (1993) and Blanchard (1993), conclude that consumption shocks or declines in consumer confidence "caused" the recession.

One might doubt that agents in the economy can forecast so much better than economists. We too are consumers, and we spend more time reading the paper and poring over the data than most. But this argument forgets aggregation. Each person has information about his own prospects, most of which is idiosyncratic. Total consumption aggregates all this information about aggregate activity. Ask a consumer about next year's GDP and he will answer "I don't know." But he may know that his factory is closing, and hence he is consuming less. This idiosyncratic shock is correlated with future GDP. Summing over consumers, aggregate consumption can reveal information about future aggregate activity, although neither consumers in the economy nor economists who study it can name what the crucial pieces of information are.

Jaimovich-Rebelo AER 2009

Our model economy is populated by identical agents who maximize their lifetime utility (U) defined over sequences of consumption (C_t) and hours worked (N_t):

(1)
$$U = E_0 \sum_{t=0}^{\infty} \beta^t \frac{(C_t - \psi N_t^{\theta} X_t)^{1-\sigma} - 1}{1-\sigma},$$

where

(2)
$$X_t = C_t^{\gamma} X_{t-1}^{1-\gamma},$$

and E_0 denotes the expectation conditional on the information available at time zero. We assume that $0 < \beta < 1$, $\theta > 1$, $\psi > 0$, and $\sigma > 0$. Agents internalize the dynamics of X_t in their maximization problem. The presence of X_t makes preferences non–time-separable in consumption and hours worked. These preferences nest as special cases the two classes of utility functions most widely used in the business cycle literature. When $\gamma = 1$ we obtain preferences of the class discussed in King, Plosser, and Rebelo (1988), which we refer to as KPR. When $\gamma = 0$ we obtain the preferences proposed by Greenwood, Hercowitz, and Huffman (1988), which we refer to as GHH.

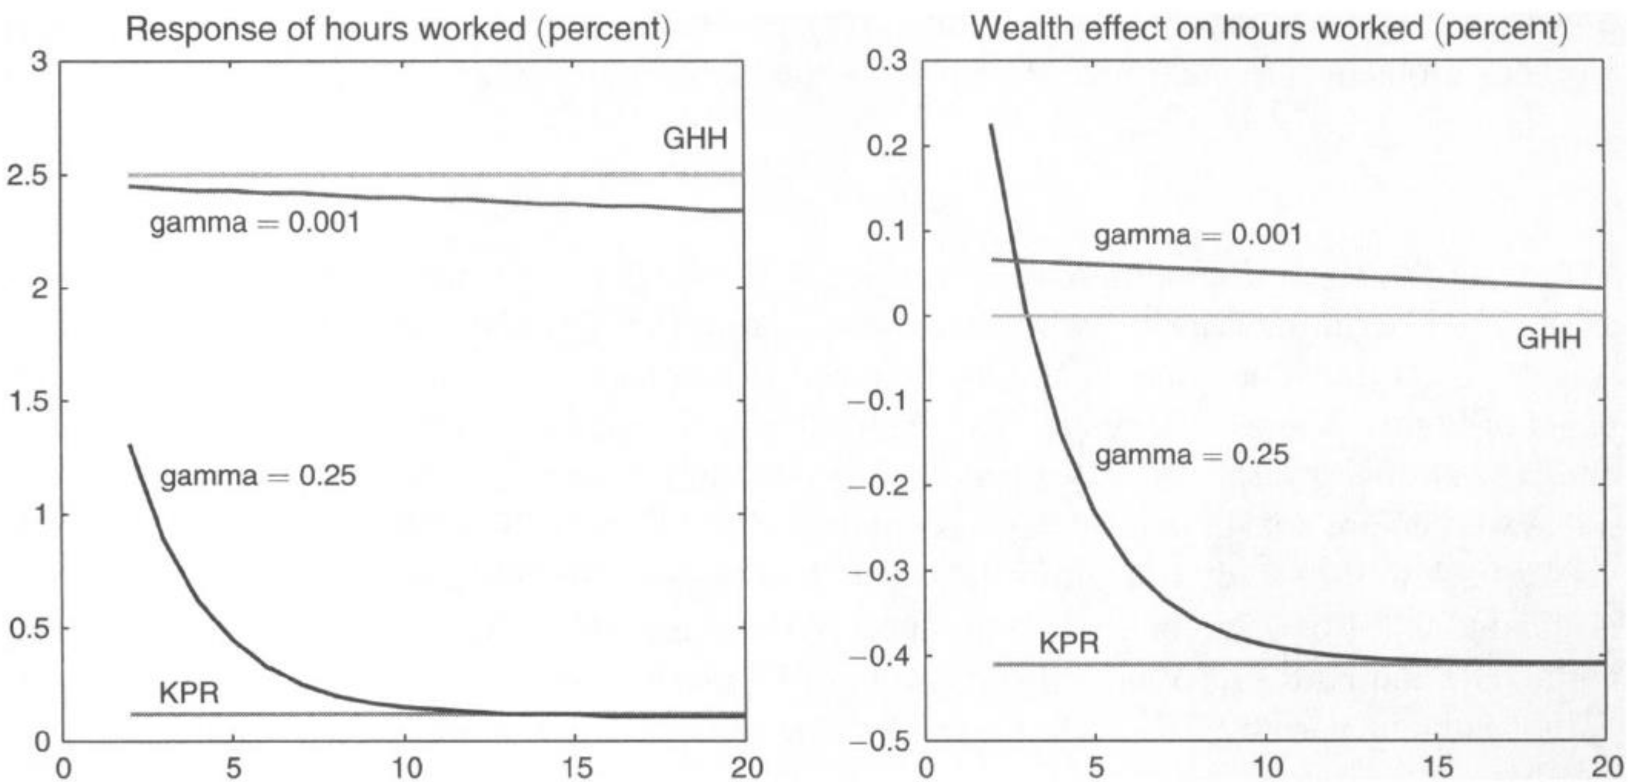

FIGURE 3. WEALTH EFFECTS ON THE LABOR SUPPLY OF A 1 PERCENT PERMANENT REAL WAGE INCREASE

We choose the following parameter values for our benchmark model. We set $\sigma = 1$, which corresponds to the case of logarithmic utility. We set θ to 1.4, which corresponds to an elasticity of labor supply of 2.5 when preferences take the GHH form. We set the discount factor β to 0.985, implying a quarterly steady-state real interest rate of 1.5 percent. The share of labor in the production function, α, is set to 0.64. We set the value of γ to 0.001, so preferences are close to a GHH specification. We choose the second derivative of the adjustment-cost functions eval-

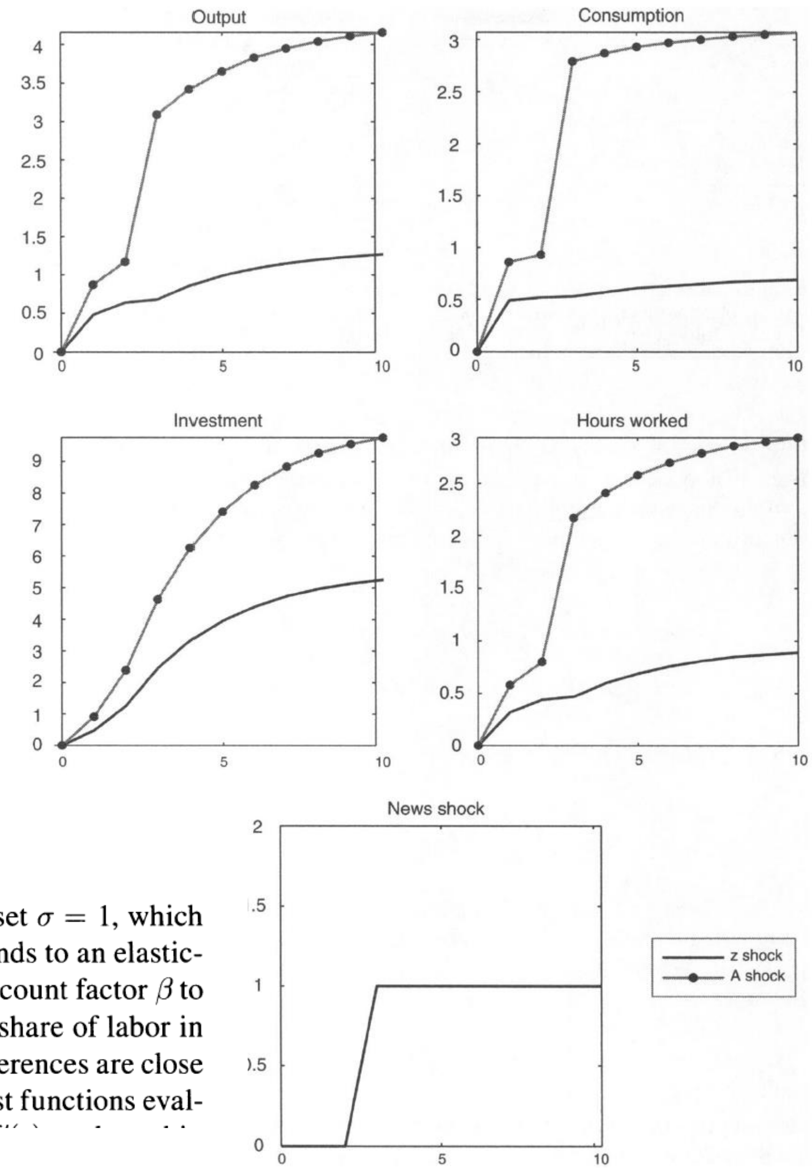

FIGURE 2. ONE-SECTOR MODEL, RESPONSE TO NEWS SHOCKS
(*Percentage deviation from steady state*)

TABLE 1—ROBUSTNESS ANALYSIS

One-sector model	News A	News z
Maximum γ	0.650	0.400
Minimum adjustment costs, $\varphi''(1)$	0.370	0.400
Minimum elasticity of labor supply $(1/(\theta-1))$	0.111	0.111
Maximum elasticity of utilization	2.500	5.000

“Information, Animal Spirits,
and the Meaning of Innovations in Consumer Confidence”
AER 2012
By Robert B. Barsky and Eric R. Sims

Overview:

- Consumer confidence (CC) measures contain a lot of information about future economic activity.

- Paper explores 2 possible explanations:

    - Animal Spirits View:

      Autonomous fluctuations in beliefs, which show up in confidence measures, have causal effects on economic activity (e.g Hall (1993), Blanchard (1993)).

    - Information or News View:

      Confidence measures contain fundamental information or news about current and future states of the economy (e.g. Cochrane (1994)).

I. Confidence and Forecasts of Economic Activity

- Measure they use the most is "E5Y." Records the responses to the following question:

"Turning to economic conditions in the country as a whole, do you expect that over the next five years we will have mostly good times, or periods of widespread unemployment and depression, or what?"

- Variable is constructed as the percentage giving a favorable answer minus the percentage giving an unfavorable answer plus 100.

- So 100 is neutral. 140 would mean fraction reflecting optimism > fraction reflecting pessimism by 40 percentage points.

- They prefer this measure eause of the relatively long horizon. Also look at other measures.

Figure 1. E5Y and E12M

Three-variable VAR

- They start with a 3-variable VAR with real GDP, real consumption expenditures, and E5Y. They use 4 lags, 1960:I – 2008:IV.

- E5Y ordered first initially.

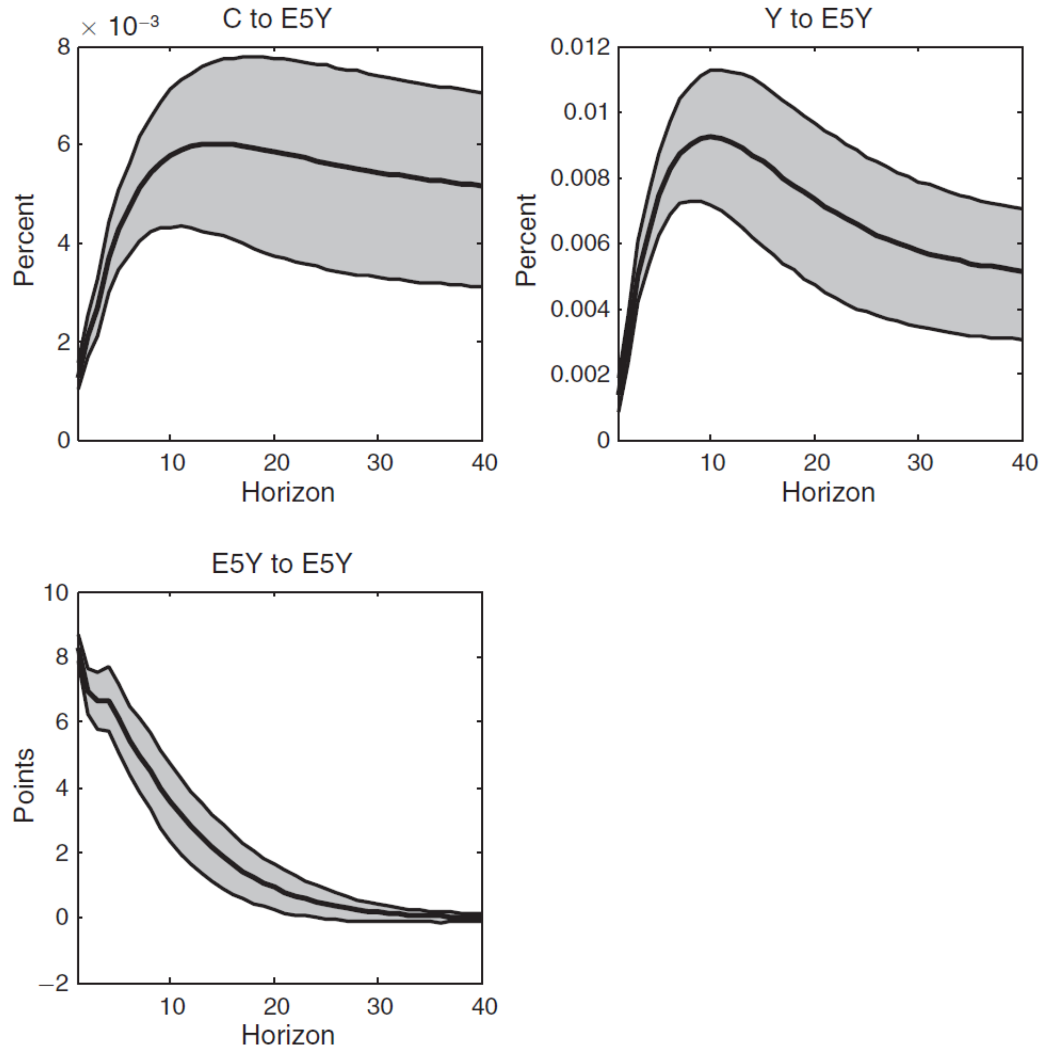

FIGURE 2. RESPONSES TO E5Y INNOVATION (*ordered first*)

Notes: These are IRFs from a three-variable VAR with E5Y, consumption, and GDP. E5Y is ordered first. The shaded areas are one-standard-error confidence bands.

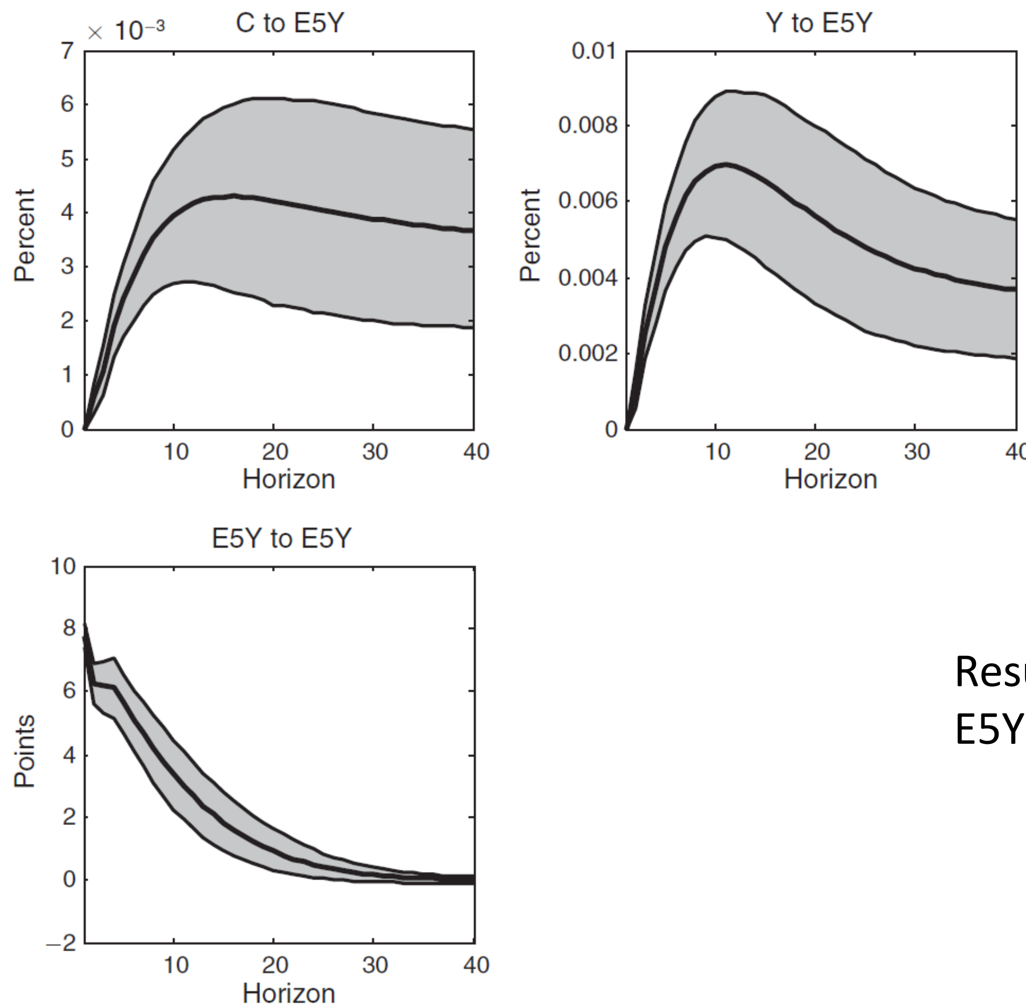

Results similar if
E5Y is ordered last.

FIGURE 3. RESPONSES TO E5Y INNOVATION (*ordered last*)

Notes: These are IRFs from a three-variable VAR with E5Y, consumption, and GDP. E5Y is ordered last. The shaded areas are one standard error confidence bands.

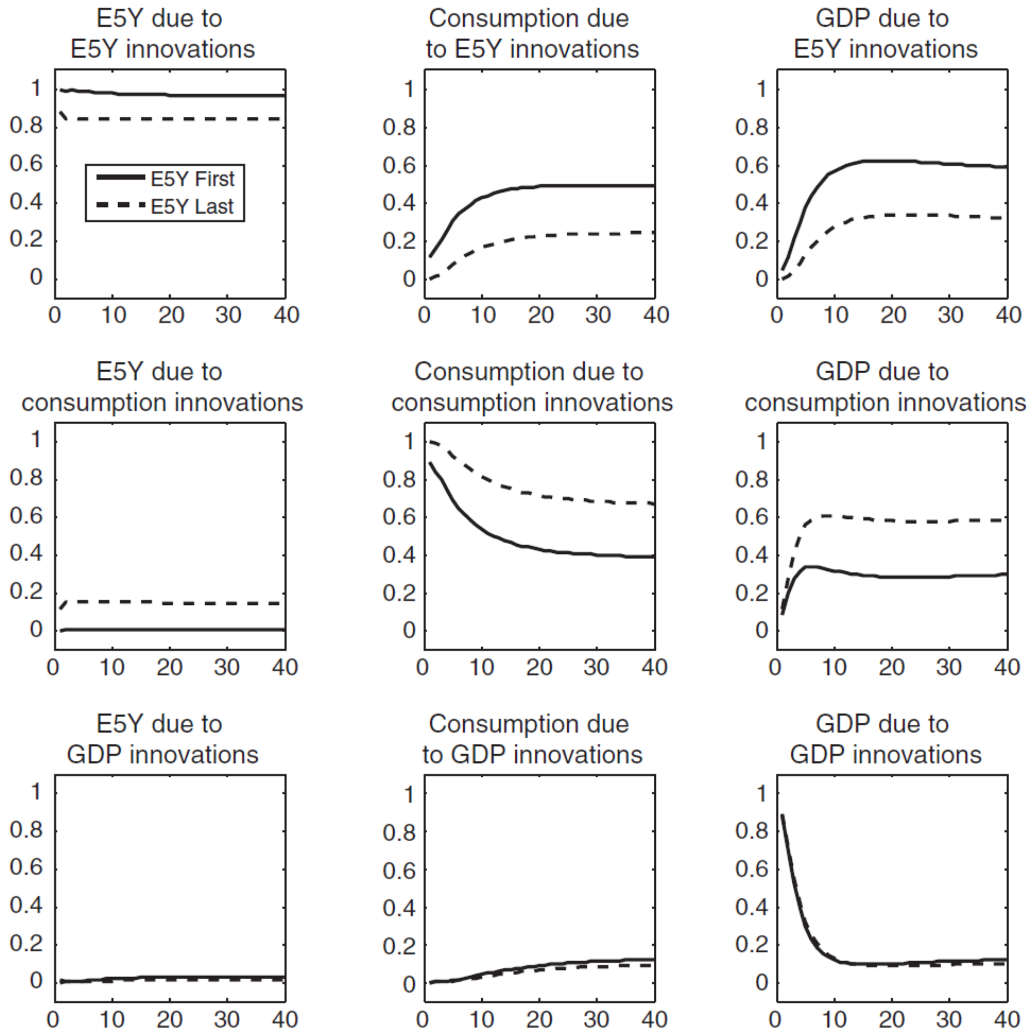

Figure 4. Variance Decomposition

Note: This figure plots variance decompositions from the three-variable VAR whose impulse responses are shown in Figures 2 and 3 under both orderings.

Innovations to confidence are an important part of forecast error variances.

II. News and Animal Spirits in a DSGE Model

- Develop a structural model to help understand the reduced-form impulse response analysis of Section I.

- Use a medium scale DSGE model. It includes price rigidity, habit formation, and adjustment costs, but also “news shocks” about future productivity growth. But agents only observe a noisy signal of this. Interpret noise innovations as animal spirits shocks.

- The following shows the key new parts of the model:

We assume that the log of neutral technology, a_t, follows a random walk with drift:

(1) $$a_t = a_{t-1} + g_{t-1} + \varepsilon_{a,t},$$

(2) $$g_t = (1 - \rho_a) g^* + \rho_a g_{t-1} + \varepsilon_{g_a,t}.$$

We assume that the drift term itself follows a stationary AR(1) process, with unconditional mean g^*. The drift term is dated $t - 1$, so that there is some predictability of technology growth. Because of this predictability, we can interpret shocks to the expected growth rate (i.e., $\varepsilon_{g_a,t}$) as "news shocks" in the sense defined by Beaudry and Portier (2004) and others. The shock $\varepsilon_{a,t}$ is the conventional surprise technology shock.

While we assume that agents can observe the level of technology period by period, we allow them to observe only a noisy signal of the growth rate. Formally:

(3) $$s_t = g_t + \varepsilon_{s,t}.$$

The shock $\varepsilon_{s,t}$ is assumed to be white noise. We will interpret it as the animal spirits shock. Following a positive animal spirits shock, the agents in the economy will erroneously expect higher subsequent productivity growth.

We assume that agents use the Kalman filter to form forecasts of the unobserved growth rate. To illustrate the mechanisms at work, Figure 6 shows impulse responses of a_t, g_t, and $g_{t|t}$ to each of the three shocks involving technology for the parameterization: $\rho = 0.8$, $\sigma_{\varepsilon_a} = 1$, $\sigma_{\varepsilon_g} = 0.1$, and $\sigma_{\varepsilon_s} = 0.1$. Note that in response to a surprise technology shock, $\varepsilon_{a,t}$, the perceived growth rate increases very slightly because agents attach some weight to the possibility that trend technology growth is on the high side but was buried in noise in the past.

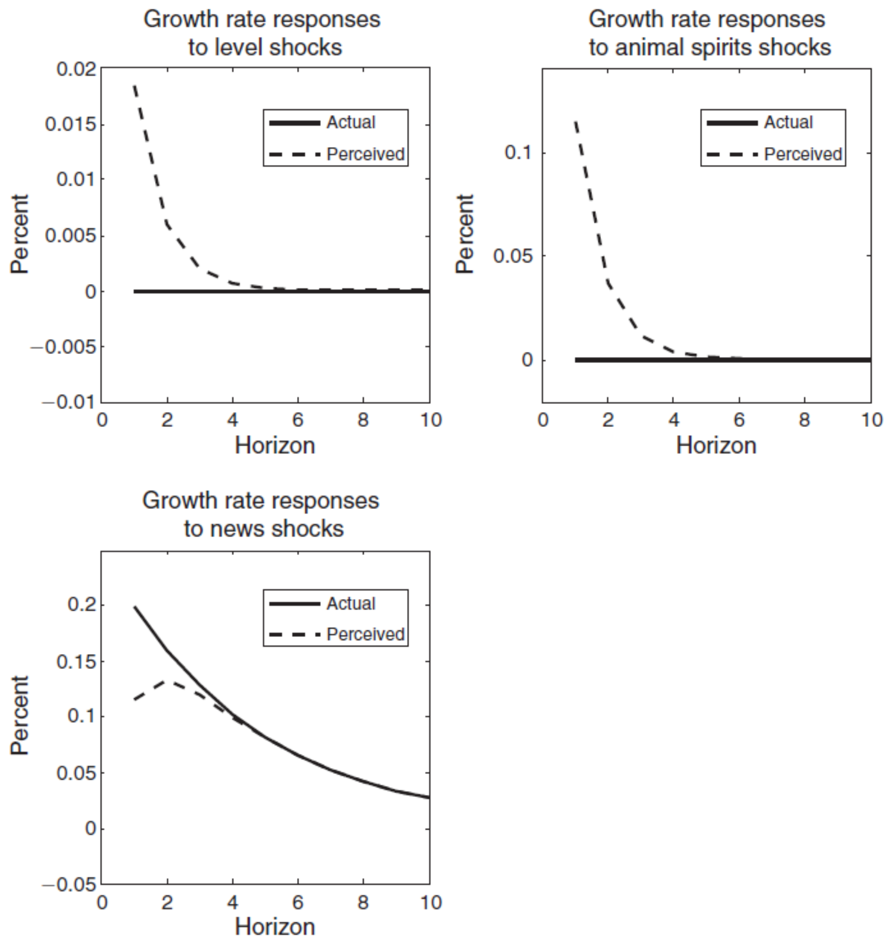

Figure 6. Actual and Perceived Growth Rates

Note: These are IRFs of true and perceived variables to various structural shocks.

growth rate. To illustrate the mechanisms at work, Figure 6 shows impulse responses of a_t, g_t, and $g_{t|t}$ to each of the three shocks involving technology for the parameterization: $\rho = 0.8$, $\sigma_{\varepsilon_a} = 1$, $\sigma_{\varepsilon_g} = 0.1$, and $\sigma_{\varepsilon_s} = 0.1$. Note that in response to a surprise technology shock, $\varepsilon_{a,t}$, the perceived growth rate increases very slightly because agents attach some weight to the possibility that trend technology growth is on the high side but was buried in noise in the past.

We assume that confidence follows a univariate first-order autoregression:

$$E5Y_t = (1 - \rho_e)E5Y^* + \rho_e E5Y_{t-1} + u_t, \tag{4}$$

where the innovation in confidence, u_t, is a function of the underlying structural shocks in the economy. Because agents cannot observe the individual structural shocks, u_t is a linear combination of the perceived innovation in the level of current technology, the perceived innovation in the expected growth rate of technology, and a pure noise term:

$$u_t = \zeta_1(a_t - a_{t-1} - g_{t-1|t-1}) + \zeta_2(g_{t|t} - \rho_a g_{t-1|t-1}) + \zeta_3 \varepsilon_{c,t}. \tag{5}$$

The shock $\varepsilon_{c,t}$ is a white noise process normalized to have variance of unity. It can be interpreted as measurement error in the confidence data.

The remainder of the model is standard and is presented in the Appendix. We

III. Estimation

- Estimate parameters by minimizing the distance between the IRFs from simulations and data.
- Calibrate the "uncontroversial parameters" and estimate the others.

For our estimation we focus on impulse responses from a somewhat larger system than that shown in Section I. In addition to confidence, consumption, and output, we also include measures of inflation and the real interest rate in the reduced-form VAR model. The reasons for this are twofold. First, because the New Keynesian model is about the interaction of real and nominal variables, the responses of inflation and interest rates help to identify the parameters of the model. Second, variation in real interest rates is a central part of the general equilibrium story when there are shocks to expectations about future technology. Our measure of inflation is the annualized

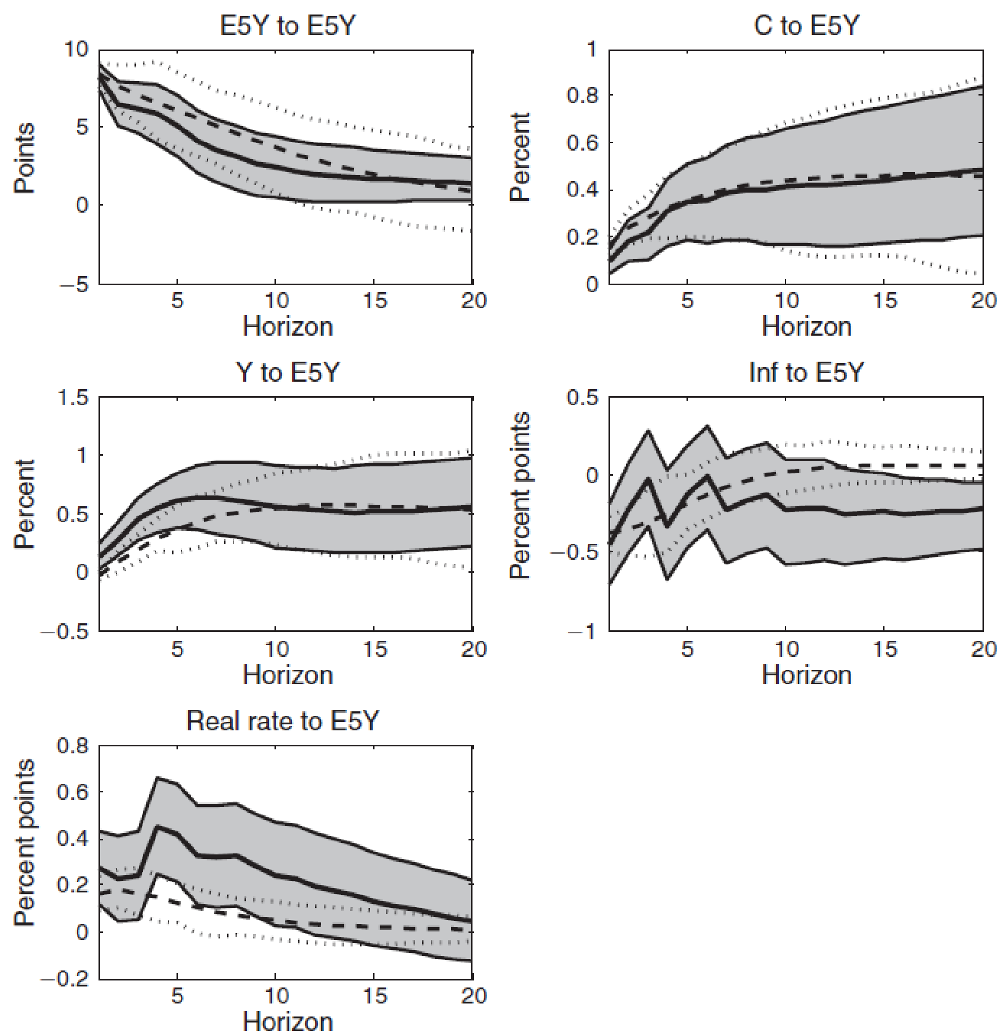

Figure 8. Simulated Responses to Confidence from Estimated Model

Notes: The solid lines are the empirical impulse responses and the shaded gray areas are the empirical confidence bands (the same as shown in Figure 7). The dashed lines are the average estimated responses from the simulations of the model at the estimated parameter values. The dotted lines are the 2.5th and 97.5th percentiles of the distribution of simulated responses.

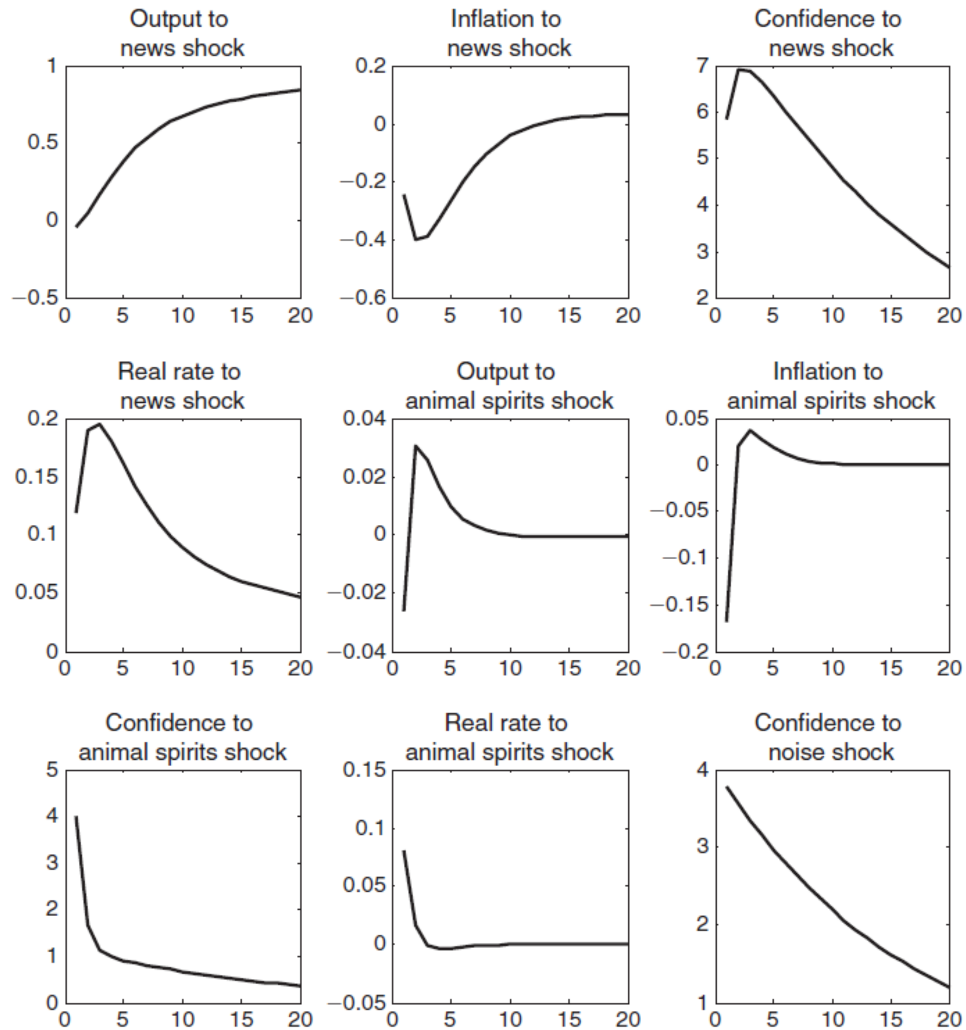

Figure 9. Theoretical Impulse Responses

Note: These are theoretical impulse responses from the model at the estimated parameter values.

TABLE 3—MODEL VARIANCE DECOMPOSITION

	$h = 1$	$h = 4$	$h = 8$	$h = 16$	$h = 20$
News					
E5Y	0.52	0.71	0.75	0.77	0.77
C	0.11	0.25	0.36	0.47	0.49
Y	0.02	0.11	0.31	0.46	0.49
Animal spirits					
E5Y	0.25	0.09	0.06	0.05	0.04
C	0.06	0.01	0.00	0.00	0.00
Y	0.01	0.01	0.00	0.00	0.00
Technology					
E5Y	0.01	0.01	0.00	0.00	0.00
C	0.43	0.48	0.50	0.48	0.47
Y	0.13	0.54	0.57	0.50	0.48
Noise					
E5Y	0.22	0.19	0.19	0.18	0.18

Note: This table shows the fraction of the forecast error variance of the respective variables explained by the structural shocks at different horizons, h, in the model of Section II at the estimated parameter values shown in Table 2.

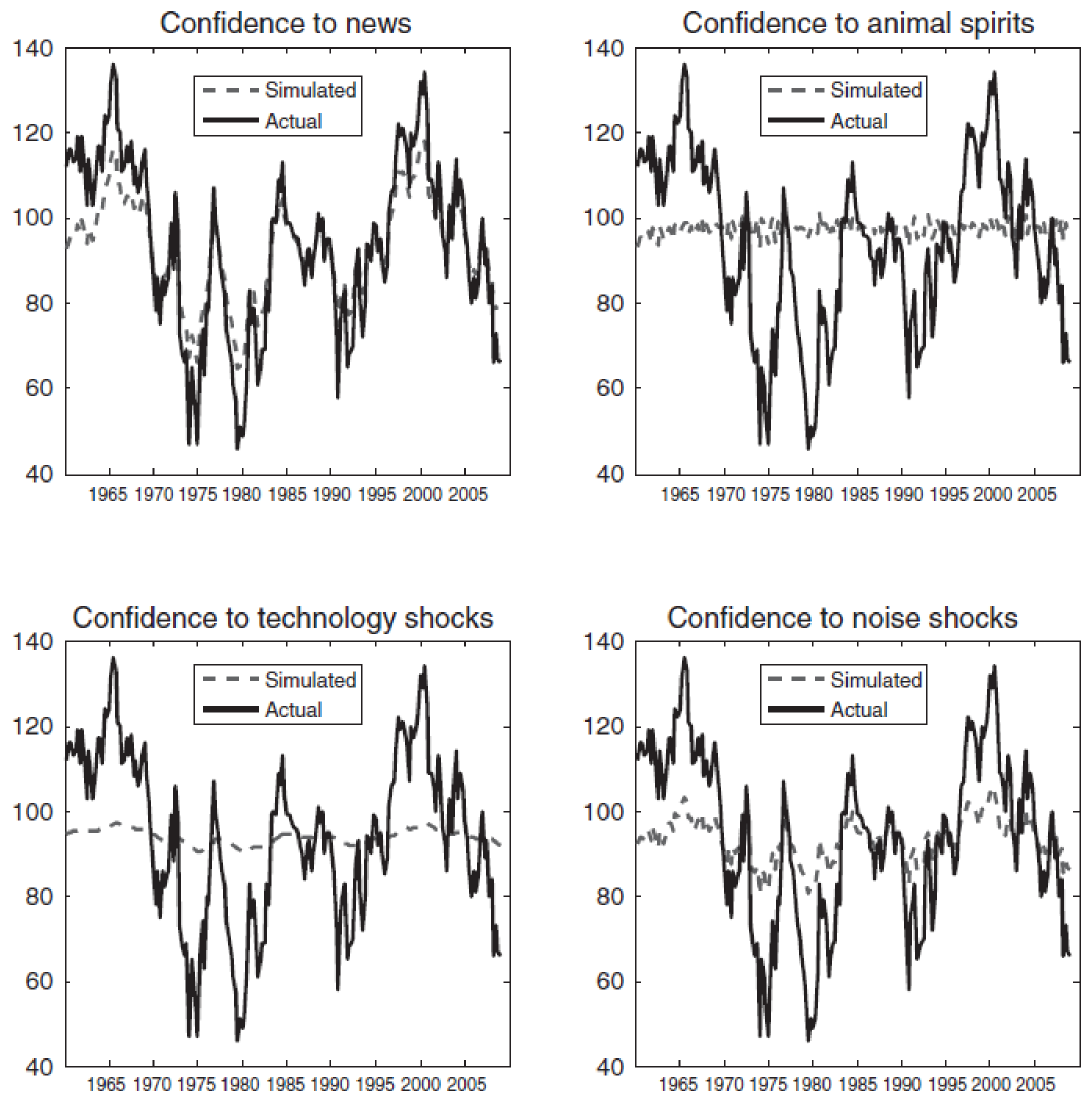

FIGURE 10. SIMULATED TIME PATHS OF CONFIDENCE

Note: The solid lines are the actual time path of E5Y; the dashed lines shows the simulated time path as if the respective shock were the only stochastic disturbance in the model.

“News Shocks and Business Cycles”
JME 2011
By Robert B. Barsky and Eric R. Sims

(Econometric details presented in class on whiteboard.)

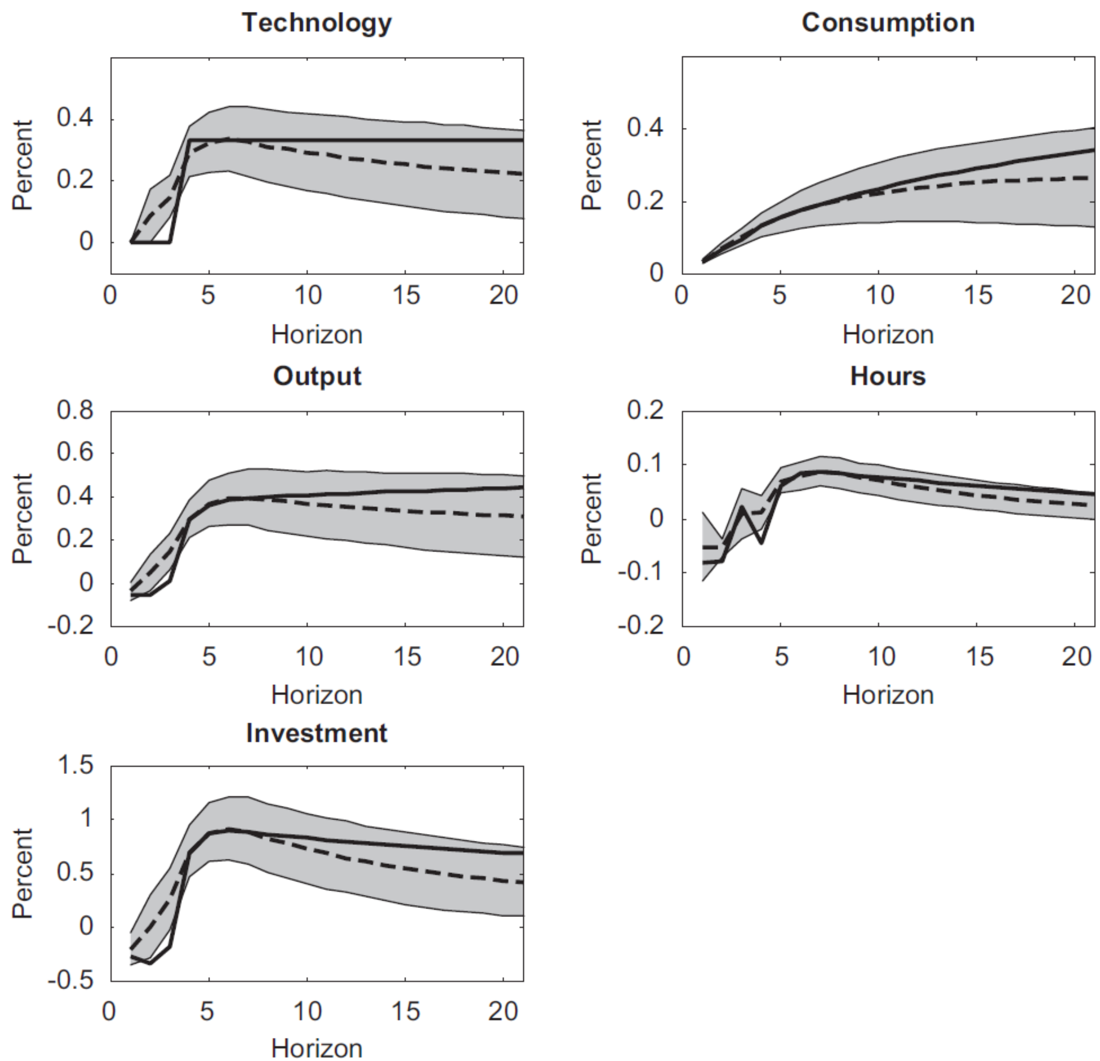

Fig. 1. Model and Monte Carlo estimated impulse responses to news shocks. The solid line shows the theoretical impulse response to a news shock from the model presented in Section 2.2. The dashed line is the average estimated impulse responses from a Monte Carlo simulation with 2000 repetitions and 191 observations per repetition. The estimated VAR includes TFP, consumption, output, and hours, all in levels. The investment response is imputed as the output response less the share-weighted consumption response. The shaded gray areas are the one $\pm$ one standard deviation confidence bands from the 2000 Monte Carlo repetitions. The horizontal axes refer to forecast horizons and the units of the vertical axes are percentage deviations (times 100).

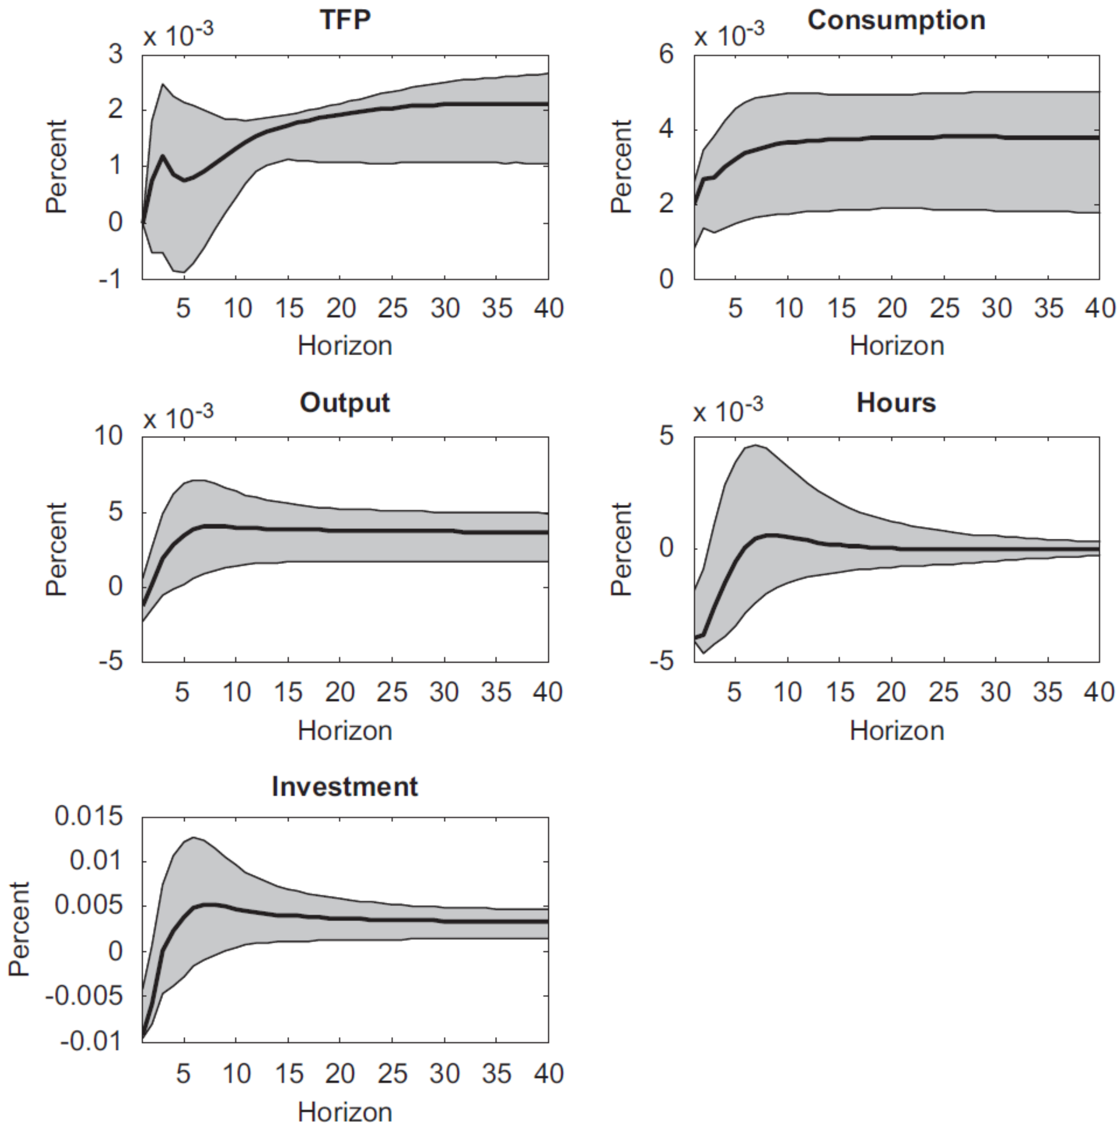

Fig. 2. Empirical impulse responses to news shock: four variable VAR. The solid lines are the estimated impulse responses to our news shock from a four variable VAR featuring TFP, consumption, output, and hours. The investment impulse response is imputed as output minus the share-weighted consumption response. The shaded gray areas are the $\pm$ one standard deviation confidence band from 2000 bias-corrected bootstrap replications of the reduced form VAR. The horizontal axes refer to forecast horizons and the units of the vertical axes are percentage deviations.

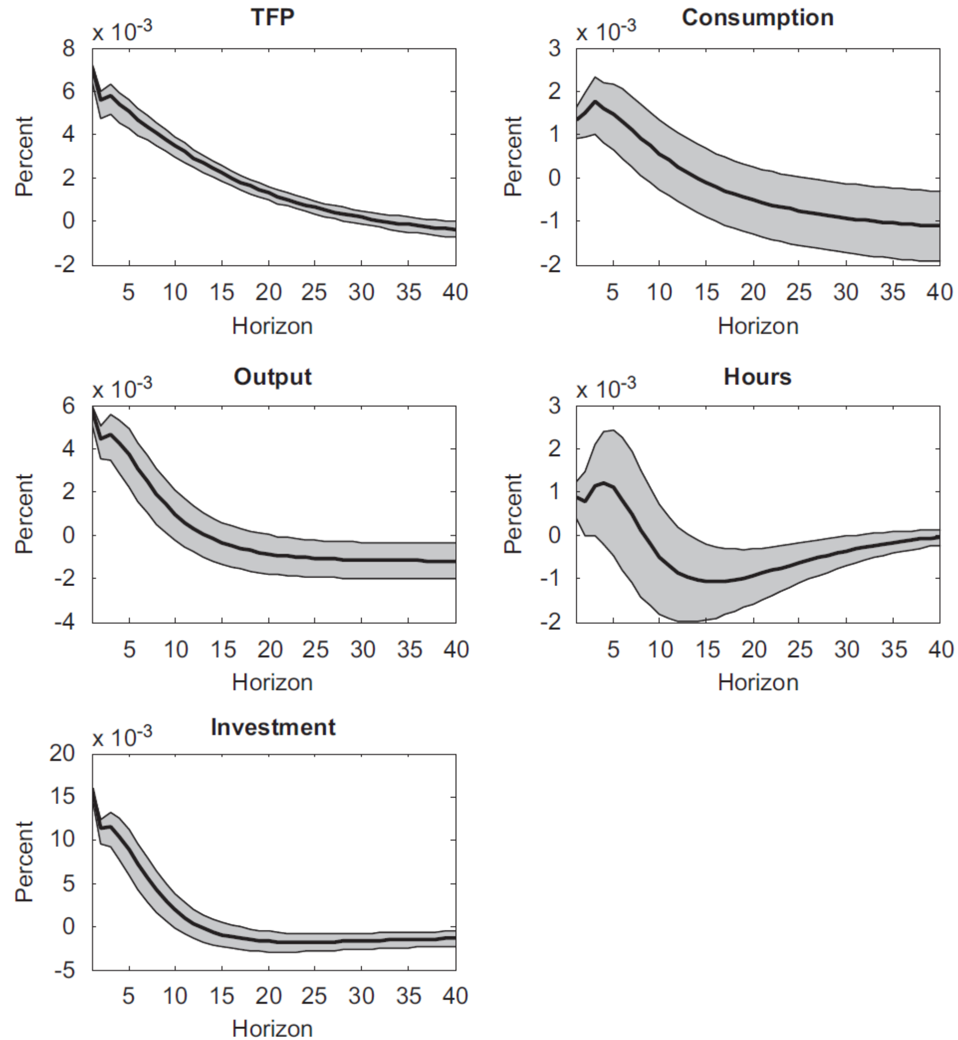

Fig. 3. Empirical impulse responses to surprise technology shock: four variable VAR. The solid lines are the estimated impulse responses to the surprise technology shock, which is simply the reduced form innovation in the VAR. The shaded gray areas are the $\pm$ one standard deviation confidence band from 2000 bias-corrected bootstrap replications of the reduced form VAR. The horizontal axes refer to forecast horizons and the units of the vertical axes are in percentage deviations.

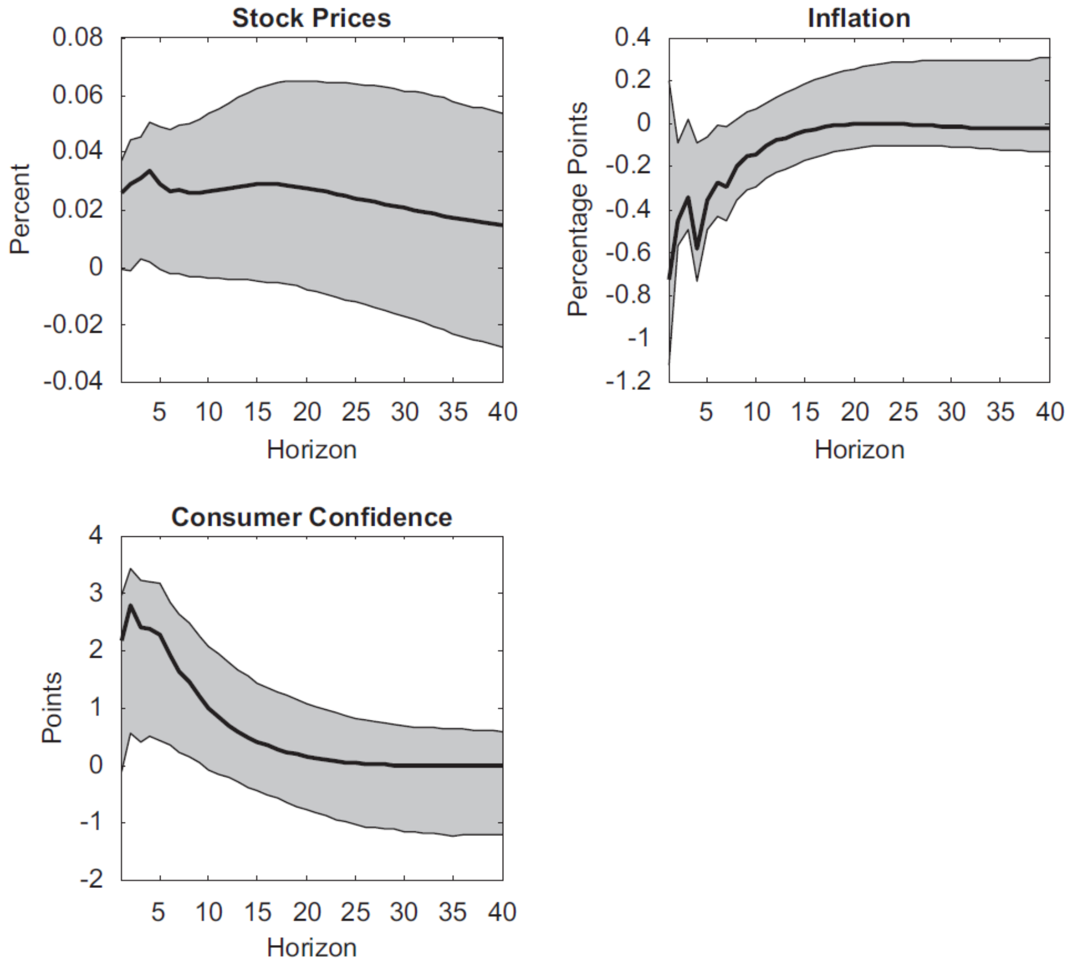

Fig. 5. Empirical impulse responses to news shock: seven variable VAR ("information variables"). These are impulse responses of the "information variables" from the seven variable VAR described in Section 3. The shaded gray areas are the $\pm$ one standard deviation confidence band from 2000 bias-corrected bootstrap replications of the reduced form VAR.

Table 1
Forecast error variance decomposition.

	$h=1$	$h=4$	$h=8$	$h=16$	$h=24$	$h=40$
TFP	0.000	0.062	0.126	0.269	0.366	0.454
	(0.00)	(0.06)	(0.11)	(0.14)	(0.15)	(0.16)
Consumption	0.050	0.234	0.377	0.493	0.524	0.507
	(0.09)	(0.18)	(0.24)	(0.27)	(0.27)	(0.26)
Output	0.111	0.091	0.242	0.382	0.429	0.431
	(0.07)	(0.10)	(0.18)	(0.23)	(0.24)	(0.24)
Hours	0.622	0.200	0.105	0.092	0.094	0.089
	(0.23)	(0.16)	(0.13)	(0.15)	(0.16)	(0.15)
Stock price	0.140	0.200	0.185	0.189	0.193	0.181
	(0.17)	(0.20)	(0.20)	(0.21)	(0.22)	(0.21)
Confidence	0.245	0.343	0.353	0.333	0.310	0.286
	(0.21)	(0.22)	(0.22)	(0.22)	(0.20)	(0.18)
Inflation	0.138	0.220	0.226	0.205	0.191	0.180
	(0.18)	(0.18)	(0.15)	(0.15)	(0.14)	(0.14)
Total TFP	1.000	0.948	0.943	0.951	0.948	0.910
Total output	0.731	0.282	0.364	0.451	0.491	0.520

The letter h refers to the forecast horizon. The numbers denote the fraction of the forecast error variance of each variable at various forecast horizons to our identified news shock. Standard errors, from a bootstrap simulation, are in parentheses. “Total TFP” shows the total variance of TFP explained by our news shock and the TFP innovation combined. “Total output” shows the total variance of output explained by the news shock and the TFP innovation combined.