News, Noise, and Fluctuations: An Empirical Exploration

Olivier J. Blanchard, Jean-Paul L'Huillier, Guido Lorenzoni *

September 2008

*MIT and NBER, MIT, and MIT and NBER. We thank Larry Christiano, Anna Mikusheva, Mark Watson, Chris Sims, and participants at the NBER Summer Institute, for very useful suggestions.

Introduction

In this paper, we explore empirically a particular class of models of fluctuations. At the core are two basic ideas: That anticipations of the future affect demand, and that demand in turn affects output in the short run.

More specifically, we think of spending decisions as depending primarily on signals about the future. These signals may be news or they may be just noise. Based on these signals, agents solve a signal extraction problem and choose spending and, because of nominal rigidities, spending affects output in the short run. If ex post, the signals turn out to have been news, agents adjust their expectations over time to the new value of the underlying fundamentals. If ex post, the signals turn out to have been noise, then the economy returns to normal over time.

We explore this class of models for two reasons. The first is that it appears to capture many of the aspects often ascribed to fluctuations, the role of animal spirits in affecting demand—"spirits" that we interpret here as coming from a rational reaction to signals about the future—, the role of demand in affecting output in the short run, together with the notion that output eventually returns to its natural level.

The second is that it appears to fit the data in a more formal way. More specifically, it offers an interpretation of structural VARs based on the assumption of two major types of shocks, shocks with permanent effects, and shocks with transitory effects on activity. As characterized by Blanchard and Quah (1989), Galí (1999), Beaudry and Portier (2006), among others, "permanent shocks" appear to lead to an increase in activity in the short run, building up to a larger effect in the long run, while—by construction—"transitory shocks" lead to a transitory effect on activity in the short run. It is tempting to associate shocks with permanent effects to news, shocks with transitory effects to noise. The interpretation of the first shock as news is made explicit in Beaudry and Portier (2006).

Our paper is organized as follows. Sections 1 and 2 present and solve our benchmark model. Section 3 looks at the use of structural VARs. It reaches a strong negative conclusion—one which came as an unhappy surprise for one of the coauthors. If the class of models we consider is a correct description of reality, then structural VARs can typically recover neither the news or noise shocks, nor their propagation mechanisms. The reason is straightforward: If agents face a signal extraction problem, and are unable to separate news from noise, then the econometrician, faced with either the same data as the agents or a subset of these data, cannot do it either. Section 4 shows however that structural estimation can be used, not to recover the time series for the news and noise shocks themselves, but to recover their variances and propagation mechanisms. It shows the results of estimation of the benchmark model. Section 5 explores a number of extensions. We find that the model fits well, and gives a clear description of fluctuations, as a result of three types of shocks: Shocks with permanent effects on productivity, with the effects on productivity slowly building over time; shocks with temporary effects on productivity, with the effects slowly decaying over time; and shocks to signals about future productivity. All three shocks affect agents' expectations, and thus affect demand and output in the short run. Over time, output adjusts to the level of productivity, thus permanently higher or lower in the case of news shocks, and back to normal in the other two cases. Section 6 concludes.

1 The model

For most of the paper, we focus on the following model, which is both analytically convenient, and, as we shall see, provides a good starting point for looking at post-war U.S. data.

We want to capture the notion that, behind productivity movements, there are two types of shocks. Shocks with permanent effects and shocks with only transitory effects. In particular, we assume that the effects of the first type of shock gradually build up over time, while the effects of the second gradually decay over time. (One can think of the transitory component as either true or reflecting measurement error. This does not matter for our purposes.)

We also want to capture the notion that spending decisions are based on agents' expectations of the future, here future productivity. We assume that agents observe productivity, but not its individual components. To capture the idea that they probably have more information than just current and past productivity, we also allow them to observe an additional signal about the permanent component of productivity. We will consider first a setup where this additional signal is not observed directly by the econometrician. Later we allow the econometrician to also observe this signal. Having solved the signal extraction problem, and based on their expectations, agents then choose spending. Because of nominal rigidities, spending determines output in the short run.

Thus, the dynamics of output are determined by three types of shocks, the two shocks to productivity, and the noise in the additional signal. For short, we shall refer to them—somewhat incorrectly but following tradition—as the "permanent shock", the "transitory shock", and the "noise shock".

Now to the specific assumptions.

1.1 Productivity

Productivity (in logs) is given by the sum of two components:

$$a_t = x_t + z_t. \tag{1}$$

The permanent component, x_t, follows a unit root process given by

$$\Delta x_t = \rho_x \Delta x_{t-1} + \epsilon_t. \tag{2}$$

The transitory component, z_t, follows a stationary process given by

$$z_t = \rho_z z_{t-1} + \eta_t. \tag{3}$$

The coefficients ρ_x and ρ_z are in $[0, 1)$, and ϵ_t and η_t are i.i.d. normal with variances σ_ϵ^2 and σ_η^2. Agents observe productivity, but not the two components separately.

In general, a given univariate representation is consistent with an infinity of decompositions between a permanent and a transitory component.[1] Here, we assume that the univariate representation of a_t is a random walk and focus on the family of processes (1)-(3) that are consistent with this assumption. We do this for two reasons. The first is analytical convenience, as it makes our arguments more transparent. The second is that, as we shall see, the assumption that the univariate representation of productivity is a random walk (with drift) provides a surprisingly good starting point when looking at post-war U.S. data. As will be clear however, our basic results, that is, the failure of SVARs and the feasibility of structural estimation, do not depend on this assumption.

The assumption that the univariate representation of productivity is a random walk conveniently restricts the family of underlying processes for x_t and z_t to a one-parameter family, parameterized by ρ. Namely, the two coefficients ρ_x and ρ_z must be equal; call their common value ρ. And the variances of the two processes must satisfy the relation $\rho\sigma_\epsilon^2 = (1-\rho)^2\,\sigma_\eta^2$. Under these restrictions, the univariate representation for productivity is given by[2]

$$a_t = a_{t-1} + u_t, \tag{4}$$

[1]See Quah (1990).

[2]This result can be proved using the fact that the spectral density of Δa_t, is equal to the sum of the spectral densities of Δx_t and Δz_t, which are, respectively, $(1-\rho e^{i\omega})^{-1}(1-\rho e^{-i\omega})^{-1}\sigma_\epsilon^2$ and $(1-e^{i\omega})(1-e^{-i\omega})(1-\rho e^{i\omega})^{-1}(1-\rho e^{-i\omega})^{-1}\sigma_\eta^2$. Under the assumed parameter restrictions this sum yields a flat spectral density.

with

$$\sigma_u^2 = \frac{1}{1-\rho^2}\sigma_\epsilon^2 + \frac{2}{1+\rho}\sigma_\eta^2.$$

Given ρ and σ_u^2, the values of σ_ϵ^2 and σ_η^2 can be computed as $\sigma_\epsilon^2 = (1-\rho)^2 \sigma_u^2$ and $\sigma_\eta^2 = \rho\sigma_u^2$. Therefore, a given random walk representation may be the result of a permanent process with small shocks that build up slowly and a transitory process with large shocks that decay slowly (a high ρ, a small σ_ϵ^2 and a large σ_η^2), or, at the other extreme, it may be the result of a permanent process which is itself close to a random walk and a transitory process close to white noise with small variance (a low ρ, a large σ_ϵ^2 and a small σ_η^2).

1.2 Consumption

We assume that consumption smoothing leads to the Euler equation

$$c_t = E[c_{t+1}|\mathcal{I}_t],$$

where $\mathcal{I}_t$ is the consumers' information at date t, to be specified below. For a generic variable X_t, we use, when convenient, $E_t[X_\tau]$ or $X_{\tau|t}$ as alternative notation for $E[X_\tau|\mathcal{I}_t]$.

We drastically simplify the supply side, by considering an economy with no capital so consumption is the only component of demand, where output is fully determined by the demand side. That is, output is given by $y_t = c_t$ and the labor input adjusts to produce y_t, given the current level of productivity. We impose the restriction that output returns to its natural level in the long run, namely that

$$\lim_{j\to\infty} E_t[c_{t+j} - a_{t+j}] = 0.$$

In Appendix A, we show that this model can be derived as the limit case of a standard New Keynesian model with Calvo pricing when the frequency of price adjustment goes to zero.

Putting the last two equations together gives

$$c_t = \lim_{j\to\infty} E_t[a_{t+j}]. \tag{5}$$

Consumption, and by implication, output, depend on the consumers' expectations of productivity in the long run.

To close the model we only need to specify the consumers' information set. Consumers observe current and past productivity, a_t. In addition, we assume that they receive a signal regarding the permanent component of the productivity process

$$s_t = x_t + \nu_t, \tag{6}$$

where ν_t is i.i.d. normal with variance σ_ν^2. We assume that the signal s_t is not observed by the econometrician. As we will see, in our benchmark model the econometrician will be able to recover s_t exactly from the data, so this assumption is not essential for our results. We assume that consumers know the structure of the model, i.e., know ρ and the variances of the three shocks.

2 Solving the model

The solution to the model gives consumption and productivity as a function of current and lagged values of the three shocks, ϵ, u, and ν. It will be convenient for later to derive it in two steps. The first, solving for consumption as a function of expectations of productivity; the second, solving for these expectations by solving the Kalman filtering problem of the consumers.

2.1 Step 1

From equations (2), (3), and (5) above,

$$c_t = x_{t,t} + \frac{\rho}{1-\rho}(x_{t|t} - x_{t-1|t})$$

or, equivalently,

$$(1-\rho)c_t = x_{t|t} - \rho x_{t-1|t} \tag{7}$$

Writing the corresponding expression for c_{t-1}, taking expectations of equation (2) at time $t-1$, and replacing, we can write consumption as

$$c_t = c_{t-1} + u_t^c, \tag{8}$$

with u_t^c given by

$$u_t^c = \frac{1}{1-\rho}\left(x_{t|t} - x_{t|t-1}\right) - \frac{\rho}{1-\rho}\left(x_{t-1|t} - x_{t-1|t-1}\right).$$

Turning to productivity, equations (1) and (3) imply

$$\begin{aligned} a_t - \rho a_{t-1} &= x_t + z_t - \rho\left(x_{t-1} + z_{t-1}\right) \\ &= x_t - \rho x_{t-1} + \eta_t. \end{aligned}$$

Adding and subtracting $x_{t|t} - \rho x_{t-1|t}$ on the right-hand side, and substituting (7) and (8), gives

$$a_t = \rho a_{t-1} + (1-\rho)\, c_{t-1} + u_t^a, \tag{9}$$

with u_t^a given by

$$u_t^a = x_t - x_{t|t-1} - \rho\left(x_{t-1} - x_{t-1|t-1}\right) + \eta_t.$$

Note that, as both current and past values of a_t and c_t are in the agents' information set at t,

$$E\left[u_t^j | a_{t-1}, c_{t-1}, a_{t-2}, c_{t-2}, ...\right] = 0,$$

for $j = c, a$. Thus, equations (8) and (9) give us the bivariate VAR representation of the joint process followed by consumption and productivity. Our assumptions imply that productivity does not help predict consumption but, if ρ is positive (and consumption and productivity are not collinear) consumption helps predict productivity.[3]

2.2 Step 2

The second step requires us to solve for the u's as a function of the underlying shocks. Agents enter the period with beliefs $x_{t|t-1}$ and $x_{t-1|t-1}$ about the current and lagged values of the permanent component of productivity. They observe current productivity $a_t = x_t + z_t$, and the signal $s_t = x_t + \nu_t$ and so update their beliefs through Kalman filtering, according to:

$$\begin{bmatrix} x_{t|t} \\ x_{t-1|t} \\ z_{t|t} \end{bmatrix} = A \begin{bmatrix} x_{t-1|t-1} \\ x_{t-2|t-1} \\ z_{t-1|t-1} \end{bmatrix} + B \begin{bmatrix} a_t \\ s_t \end{bmatrix} \tag{10}$$

where the matrices A and B depend on the underlying parameters, ρ and $\sigma_\epsilon, \sigma_\eta, \sigma_\nu$ (see Appendix B).

[3]Note that, under our assumptions, both productivity and consumption have random walk *univariate* representations. For productivity, it is simply by assumption. For consumption, it follows from equation (5), independent of the form of the process for productivity.

2.3 The dynamic effects of shocks

Equations (8), (9) and (10), together with equations (1), (2), and (3), characterize the dynamic responses of productivity and consumption to the shocks. Except in two special cases to which we shall come back below (the case of a fully informative or a fully uninformative signal), these must be solved numerically.

Figure 1 gives the computed impulse responses of consumption and productivity to the three shocks. The parameters are chosen to be roughly in line with the estimates we obtain later, in Section 4. The time unit is the quarter. The parameter ρ is chosen equal to 0.97: this implies slowly building permanent shocks and slowly decaying transitory shocks. Given that the standard deviation of the innovation to productivity, σ_u, is roughly equal to 0.8%, this, together with the value of ρ, implies standard deviations of the two technology shocks, σ_ϵ and σ_η, equal to 0.03% and 0.8%, respectively. The standard deviation of the noise shock, σ_ν, is set to 2.0%, implying a fairly noisy signal.

In response to a one-standard deviation increase in ϵ, a permanent technology shock, productivity builds up slowly over time—the implication of a high value for ρ. Consumption also increases slowly. This reflects the fact that the standard deviations of the transitory shock, η, and the noise shock, ν, are both large relative to the standard deviation of ϵ. Thus, it takes a long time for consumers to be able to assess that this is really a permanent shock and to fully adjust consumption.

For our parameter values, consumption (equivalently, output) initially increases more than productivity, generating a transitory increase in employment. Smaller transitory shocks, or a more informative signal would lead to a larger initial increase in consumption, and thus a larger initial increase in employment. Larger transitory shocks, or a less informative signal, might lead instead to an initial decrease in employment.

In response to a one-standard deviation increase in η, the transitory shock, productivity initially increases, and then slowly declines over time. As agents put some weight on it being a permanent shock, they initially increase consumption. As they learn that this was a transitory shock, consumption returns back to normal over time. For our parameter values, consumption increases less than productivity, leading to an initial decrease in employment. Again, for different parameters, the outcome may be an increase or a decrease in employment.

Finally, in response to a one-standard deviation increase in ν, the noise shock, consumption increases, and then returns to normal over time (the response of consumption need not be monotonic; in the simulation presented here, consumption turns briefly negative, before returning to normal). By assumption, productivity does not change, so employment initially

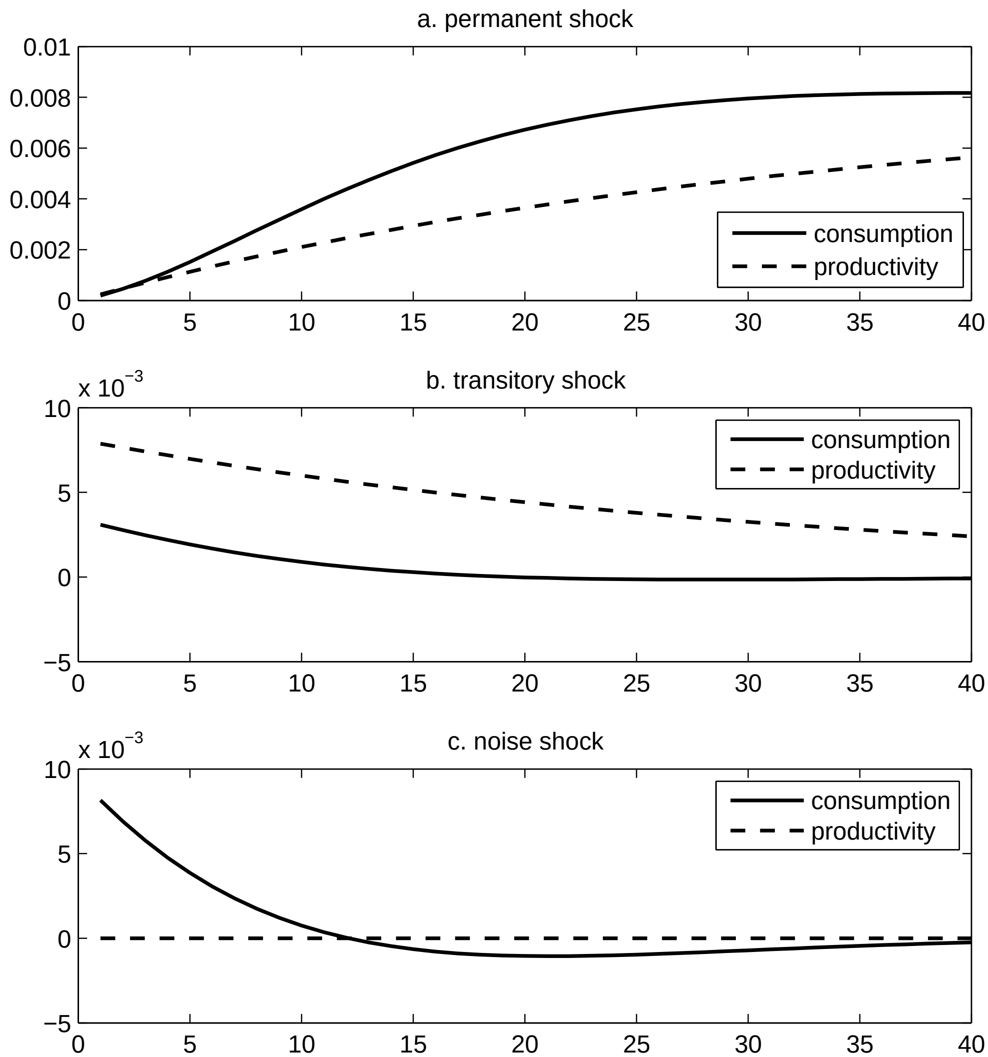

Figure 1: Impulse Responses to the Three Shocks

increases, to return to normal over time.

3 A structural VAR approach

The question we take up in this section is whether a structural VAR approach can recover the underlying shocks and their impulse responses.

Given that we (the econometrician) observe only two variables, productivity and consumption, and there are three shocks, the answer is trivially no, unless the model is degenerate. But one may still hope that the SVAR based on long run restrictions, will identify the effect of ϵ, the only shock with long run effects on productivity and consumption. As we shall see, the answer however is no: In an environment in which agents face a non-trivial signal extraction problem, SVARs will typically fail.

It is best to start with two special cases of the model we introduced earlier, and then to consider the general case.

3.1 A fully uninformative signal

Consider first the case of a fully uninformative signal, $\sigma_\nu = \infty$, so the consumers' only information is given by current and past values of a_t.

Then, trivially, our random walk assumption for a_t leads to $c_t = a_t$. In this case, the two innovations u_t^c and u_t^a coincide and are identical to the innovation u_t in the univariate representation of a_t. The bivariate dynamics of consumption and productivity are given by

$$\begin{aligned} c_t &= a_{t-1} + u_t, \\ a_t &= a_{t-1} + u_t. \end{aligned}$$

This characterization holds for any value of ρ. Thus, whatever the value of ρ and the relative persistence and importance of the permanent and transitory components of productivity, a structural VAR with long-run restrictions will attribute all movements in productivity and consumption to permanent shocks, and none to transitory shocks. The impulse responses of productivity and consumption to ϵ will show a one-time permanent increase; the impulse responses of productivity and consumption to η will be identically equal to zero.

3.2 A fully informative signal

Consider next the case of a fully informative signal, $\sigma_\nu = 0$, so consumers no longer face a signal extraction problem. They know exactly the value of the permanent component of productivity, x_t—and by implication, the value of the transitory component, $z_t = a_t - x_t$. In this case, equations (8) and (9) simplify to:

$$\begin{aligned} c_t &= c_{t-1} + \frac{1}{1-\rho}\epsilon_t, \\ a_t &= \rho a_{t-1} + (1-\rho)\, c_{t-1} + \epsilon_t + \eta_t. \end{aligned}$$

Consumption responds only to the permanent shock; productivity to both. In this case, a structural VAR approach will indeed work. Imposing the long-run restriction that only one of the shocks has a permanent effect on consumption and productivity will recover ϵ_t and η_t, and, by implication, their dynamic effects.

3.3 The general case

Which of these two cases is pathological? The answer is, unfortunately, the second. As soon as the signal is not fully informative, so consumers face a signal extraction problem, the structural VAR approach will fail. Figure 2 shows the estimated IRFs to the shocks with permanent and transitory effects obtained from structural VAR estimation, together with the true IRFs to the three underlying shocks. The underlying parameters are chosen to be the same as for Figure 1. (The estimated IRFs are obtained by generating long time series for consumption and productivity using the true model, and then running a structural VAR using these time series).

Look first at the true and estimated IRFs of productivity to a permanent shock. The black line (red in pdf) in the top left quadrant replicates the corresponding IRF in Figure 1, namely a small initial effect, followed by a steady buildup over time. The dotted line gives the estimated IRF from SVAR estimation: The initial effect is much larger, the later buildup much smaller. Indeed, simulations show that the less informative the signal, the larger the estimated initial effect, the smaller the later build up. (Remember that, when the signal is fully uninformative, the estimated IRF shows a one-time increase, with no further build up over time).

Turn to the true and estimated IRFs of consumption to a permanent shock in the bottom left quadrant. The black line (red in pdf) again replicates the corresponding IRF in Figure 1, showing a slow build-up of consumption over time. The dotted line shows the estimated

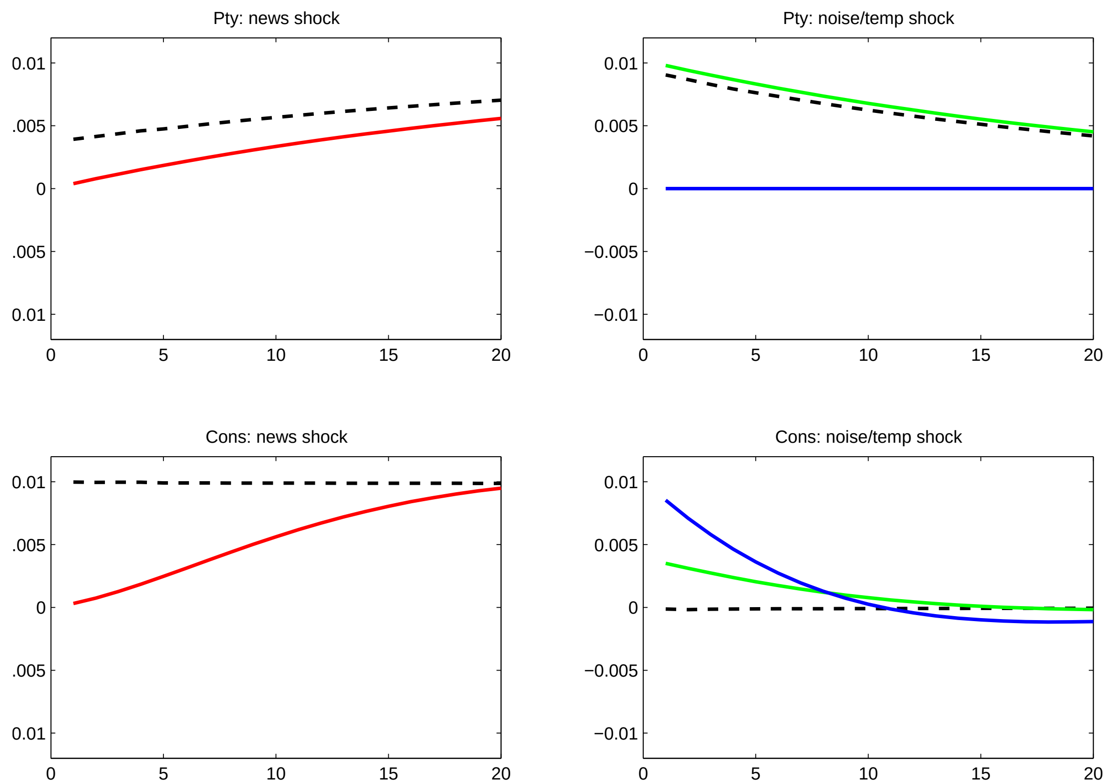

Figure 2: True and SVAR-based estimated IRFs

IRF, namely a one-time response of consumption with no further build up over time.

The right quadrants show the true and estimated responses to transitory shocks. The grey lines (green in pdf) show the IRFs to a transitory shock, the black lines (blue in pdf) show the IRFs to a noise shock. The dotted lines give the estimated IRFs from SVAR estimation. They show that the estimated IRF of productivity to a transitory shock is close to the true IRF to a transitory shock, and that the estimated IRF of consumption is equal to zero.

The estimated response of consumption to estimated permanent shocks (full initial response) and to estimated transitory shocks (no response) are particularly striking, and suggest a more general proposition: That, within the structure of our model (in which consumption has a univariate random walk representation), the estimated response of consumption will have these two characteristics, independent of the parameters of the model, and in particular independent of ρ and σ_ν. This proposition is indeed true, and is proven in Appendix C.

In short, the IRFs from an SVAR overstate the initial response of productivity and consumption to permanent shocks, and thus give too much weight to these shocks in accounting for fluctuations. For productivity, the less informative the signal, the larger the overstatement. For consumption, the overstatement is independent of the informativeness of the signal.

Why do SVARs fail? Because, if the consumers are confused about whether a change in productivity reflects a shock with permanent effects or a shock with transitory effects, and thus face a non-trivial signal extraction problem, the econometrician, who typically has access to even less data than the agents, faces the same problem. Put another way, if we (the econometrician), could recover the true shocks, then the consumers would have been able to do the same, and would not face a signal extraction problem in the first place. This suggests that existing results, mentioned in the introduction, give too much weight to permanent shocks in fluctuations.

3.4 What if the econometrician has more information than the agents?

The argument above suggests two potential ways out, both based on the possibility that the econometrician may have access to more information about time t when doing the estimation than agents had at the time.

The first is that, if we think of the transitory component as reflecting in part measurement error, and if the series for productivity is revised over time, the econometrician, who has access to the revised series, may be better able than the consumers to separate the permanent and the transitory components. To take an extreme case, if the transitory component reflects only measurement error, and if the revised series remove the measurement error, then the econometrician has access to the time series for the permanent component directly, and can therefore separate the two components. While this is extreme, this suggests that the bias from SVAR estimation may be reduced when using revised series rather than originally published series. [4]

The second starts from the observation that the econometrician, when doing estimation, observes realizations of productivity after time t. Thus, while the agents can only do Kalman filtering, the econometrician can do Kalman smoothing, that is use the information available after time t to get estimates of ϵ_t and η_t. This however fails as well. Take, for simplicity, the case where the signal is fully uninformative so there is no additional information in observing consumption, and we can assume the econometrician just observes productivity.

In this case, the econometrician can, by estimating the univariate representation of productivity, get estimates of u_t. Can he use future values of u_t to infer the values of ϵ_t and η_t at time t? The answer is no, and follows from our earlier indeterminacy result. Recall that the univariate process is consistent with an infinity of underlying permanent and transitory components, parameterized by ρ. There is no way to learn ρ from the data, and thus to get estimates of ϵ_t and η_t. We will return to this theme when we discuss how, using current and future data, the econometrician can recover the current shocks after structural estimation of the model. [5]

4 Structural estimation

We now turn to structural estimation, proceeding in two steps. For the benchmark model we have introduced, structural estimation is particularly easy, and all parameters (save one)

[4] A related article here is Rodriguez Mora and Schulstad (2007). They show that growth in period t is correlated with preliminary estimates of past growth available in period t, not with final estimates, available later. One potential interpretation of these results is that agents choose spending in response to these preliminary estimates, and their spending in turn determines current output.

[5] Suppose the econometrician knew ρ, used Kalman smoothing to get estimates of the underlying shocks, and then used these estimated shocks to get IRFs of productivity and consumption to these shocks. Would he get the right IRFs? The answer is still no. Even with an infinite amount of data, the econometrician cannot recover the exact values of ϵ and η. For example, the estimated time series for ϵ from Kalman smoothing gives a series with high serial correlation—despite the fact that the true ϵ is i.i.d.

can be obtained using OLS; thus we start with it and show the results. For more general processes however, one must use maximum likelihood. We show how it can be done, and show estimation results.

4.1 A simple OLS approach

It is clear that we (the econometrician) cannot recover the three shocks from the two variables we observe, productivity and consumption. It is clear however that, from estimating equation (9) we can recover ρ, and from estimating equation (4), we can recover σ_u. Given ρ and σ_u, we can recover σ_ϵ and σ_η. Given those, we can then use the model to characterize the dynamic effects of permanent and transitory shocks. Recovering the variance of the noise shock is less straightforward, but it can be done looking at other estimated moments (in particular, the covariance between u_t^c and u_t^a is monotone increasing in σ_ν).

How well does our simple benchmark model fit the time series facts for productivity and consumption? The answer is: fairly well. Although it clearly misses some of the dynamics in the data, it seems worth starting with it.

The basic characteristics of the two time series are shown in Table 1. We construct the productivity variable as the logarithm of the ratio of GDP to employment. We construct the consumption variable as the logarithm of the ratio of NIPA consumption to population. We use quarterly data, from 1970:1 to 2008:1. An issue we have to confront is that, in contradiction to our model, and indeed to any balanced growth model, the productivity and consumption variables have different growth rates over the sample (0.34% per quarter for productivity, versus 0.48% for consumption). This difference reflects factors we have left out of the model, from changes in participation, to changes in the saving rate, to changes in the capital-output ratio. For this reason, we run the two variables on linear time trends, and use the residuals in what follows.[6]

Lines 1 and 2 of Table 1 show the results of estimated AR(1) for the first differences of the two variables. Recall that our model implies that both productivity and consumption should follow random walks, so the AR(1) term should be equal to zero. In both cases, the AR(1) term is indeed small, insignificant in the case of productivity, significant in the case of consumption.

Our model further implies a simple dynamic relation between productivity and consump-

[6]We are aware that, in the context of our approach, where we are trying to isolate potentially low frequency movements in productivity, this is a rough and dangerous approximation. But, given our purposes, it seems to be a reasonable first pass assumption.

tion. Rewriting equation (9) as a cointegrating regression gives:

$$\Delta a_t = (1-\rho)(c_{t-1} - a_{t-1}) + u_t^a$$

Line 3 shows the results of estimation of this equation. Line 4 allows for lagged rates of change of consumption and productivity, and shows the presence of richer dynamics than implied by our specification, with small but significant coefficients on lagged rates of change consumption and productivity.

Line	Dependent variable:	$\Delta a(-1)$	$\Delta c(-1)$	$(c-a)(-1)$
1	Δa	-0.06 (0.7)		
2	Δc		0.23 (3.0)	
3	Δa			0.05 (1.2)
4	Δa	-0.21 (-2.3)	0.32 (3.4)	0.02 (0.4)

Table 1: Consumption and Productivity Regressions.
Sample: 1970:1 to 2008:1. t-statistics in parentheses.

The estimation shown in line 3 implies a value of ρ of 0.95. Together with an estimated standard deviation for σ_u of 0.7%, these imply $\sigma_\epsilon = 0.035\%$ and $\sigma_\eta = 0.7\%$. In words: These results imply a very smooth permanent component, in which small shocks steadily build up over time, and a large transitory component, which decays slowly over time.

The fact that we are able in our benchmark model to recover the central parameter ρ from a simple OLS regression is clearly a special case. For more general specifications of productivity or consumption behavior, one must adopt a different approach. We now show this general approach, and then return to the data.

4.2 Kalman Filtering and Maximum Likelihood

The estimate a model where consumers face a non trivial signal extraction problem, one can, generally, proceed in two steps.

- Take the point of view of the consumers. Write down the dynamics of the unobserved states in state space representation and solve the consumers' filtering problem. In our case, the relevant state for the consumer is given by $\xi_t \equiv (x_t, x_{t-1}, z_t)$, its dynamics are given by (2) and (3), the observation equations are (1) and (6), and Kalman filtering gives us the updating equation (10).

- Next, take the point of view of the econometrician, and, again, write down the model dynamics in state space representation and write the appropriate observation equations (which depend on the data available). In our case, the relevant state for the econometrician is given by $\xi_t^E \equiv (x_t, x_{t-1}, z_t, x_{t|t}, x_{t-1|t}, z_{t|t})$. Notice that the consumers' expectations become part of the unobservable state and the consumers' updating equation (10) becomes part of the description of the state's dynamics. The observation equations for the econometrician are now (1) and (7), where the second links consumption (observed by the econometrician), to the consumer's expectations. The econometrician's Kalman filter can then be used to construct the likelihood function and estimate the model's parameters.

Table 2 shows the results of estimation of the benchmark model, presented as a grid over values of ρ from 0.0 to 1.00. The maximum likelihood is reached for $\rho = .96$, values of σ_ϵ and σ_η of 0.03% and 0.8% respectively. The standard deviation of the noise, σ_ν is 0.2% (although, as one can see from the table, this standard deviation is very sensitive to small changes in ρ.)

Line	ρ	σ_u	σ_ϵ	σ_η	σ_ν	ML
1	0.00	0.83	0.83	0.00	0.80	545.3
2	0.20	8.19	6.55	3.66	0.00	664.3
3	0.40	4.36	2.62	2.76	0.00	868.7
4	0.60	2.42	0.97	1.88	0.07	1088.8
5	0.80	1.25	0.25	1.12	0.01	1355.6
6	0.90	0.92	0.09	0.87	0.00	1481.6
7	0.95	0.82	0.04	0.80	0.00	1521
8	0.96	0.82	0.03	0.80	0.19	1527.2
9	0.97	0.85	0.03	0.84	2.18	1524.2
10	0.98	1.57	0.03	1.55	1.00	1398.7
11	0.99	1.48	0.01	1.47	4.01	1386.3
12	1.00	0.87	0.00	0.87	0.00	1382.3

Table 2: Maximum Likelihood Estimation of the Benchmark Model

One simple exercise, using this approach, is to relax the random walk assumption for productivity, allowing ρ_x to differ from ρ_z, and allowing the variances of the shocks to be freely estimated. The results of estimation are $\rho_x = 0.98$, $\rho_z = 0.96$, $\sigma_e psilon = .02\%$, $\sigma_\eta = 0.8\%$, and $\sigma_\nu = 1.9\%$. Thus, except for the standard deviation of the signal, the results are very close to those obtained under the random walk assumption.

What do our results imply, in terms of dynamic effects of the shocks, and variance decompositions. If we use the estimated parameters from the benchmark model (line 8 in Table 2), the dynamic effects of each shock were given in Figure 1 earlier, and we discussed them already: A slow and steady build up of permanent shocks on productivity and consumption; a slowly decreasing effect of transitory shocks on productivity and consumption; and a slowly decreasing effect of noise shocks on consumption.

Another way of looking at the results is in terms of variance decompositions. These are given in Table 3. They show that short run fluctuations are dominated by transitory and noise shocks, rather than by permanent shocks.

[...]

Table 3: Variance Decomposition

4.3 Recovering shocks and states

So far we have focused on using structural estimation to estimate the model's parameters. We can then address the question: what information about the actual realizations of the shocks ϵ_t, η_t, and ν_t, can be recovered from this structural estimation? Recovering the shocks also allows us to form estimates of the unobservable states x_t and z_t, so as to reconstruct a retrospective story of observed cycles.

We know from our discussion of the SVAR approach, that the information in current and past values of c_t and a_t is not sufficient to derive the values of the current shocks. In fact, we will see that even if the econometrician had an infinite series of future data on c_t and a_t, it would typically not be possible to exactly recover ϵ_t, η_t, and ν_t. However, this does not mean that the data contain no information on these shocks. In particular, using the Kalman smoother the econometrician can update its beliefs on ϵ_t, η_t, and ν_t at t,$t+1$, etc. and achieve smaller and smaller mean squared errors. In Figure 3 we plot the mean squared errors of the smoothed estimates of the shocks, using the parameter estimates for our benchmark model. The figure displays mean squared errors of the following form

$$E_{t+j}[(\epsilon_t - E_{t+j}[\epsilon_t])^2],$$

for $j = 0, 1, 2, ...$ (expectations are taken using the econometrician's information sets).[7] Each

[7]The mean squared error are computed in the steady state of the Kalman filter, that is, assuming the

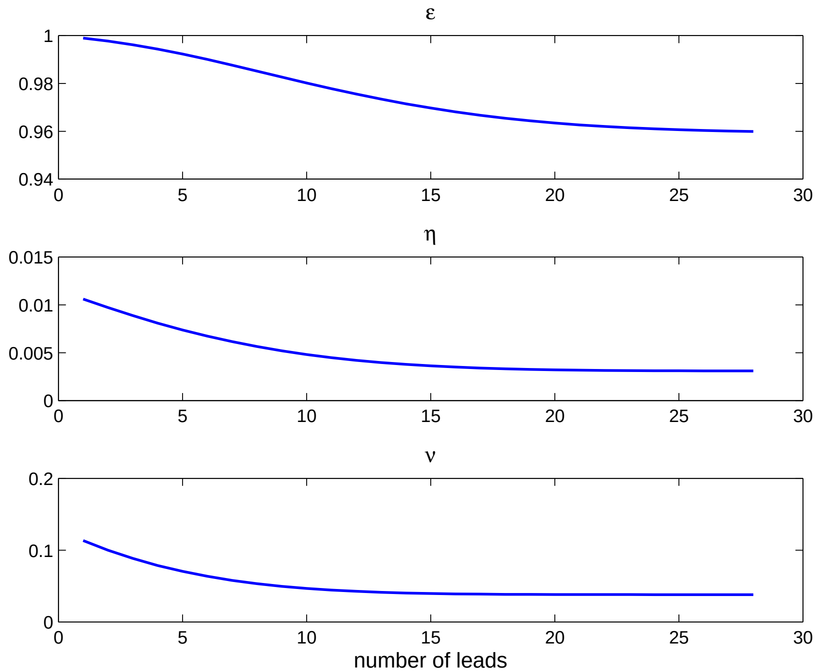

Figure 3: Normalized MSE of the estimated shocks at time t using data up to $t+j$

MSE is normalized using the ex ante variance of the respective shock, that is, σ_ϵ^2, σ_η^2, and σ_ν^2. Notice that the transitory shock and the noise shock η and ν can be quite precisely estimated already on impact and the precision of their estimates almost doubles in the long run. On the other hand, the permanent shock ϵ is considerably harder to estimate, and even when infinite future data are available, the residual variance is about 96% of the prior uncertainty on the shock.

Figure 4 presents a similar exercise for the underlying states x and z.[8]

5 Extensions

We have shown how models where agents face signal-extraction problems cannot be estimated through SVARs, but can be estimated through structural estimation. Structural estimation however requires a full specification of the model, including the processes for the permanent and transitory components of productivity, the information structure, the behavior of consumers. And, unfortunately, the estimated parameters are likely to be sensitive to

econometrician has a very long (infinite) series of past data but only $t+j$ future data.

[8]This time the MSE are not normalized as the ex ante variance of x_t would be infinity.

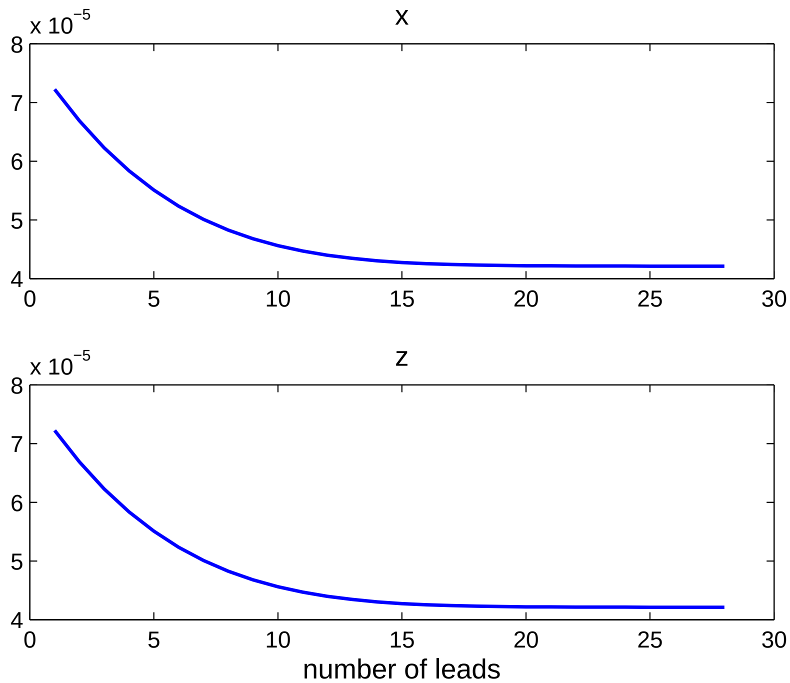

Figure 4: MSE of the estimated states at time t using data up to $t+j$

the specific assumptions.

There are at least two dimensions in which we think our benchmark model needs to be extended.

The first is motivated by the data. As we saw from Table 1, the dynamics of consumption, and the dynamic relation between productivity and consumption, are richer than those implied by the benchmark. These require at least a modification of our assumptions about consumption behavior. Our assumption about consumption implies that consumption follow a random walk for any productivity process and any standard deviation of the noise in the signal. As we have seen however, the univariate process for consumption, shown in line 2 of Table 2, shows evidence of richer dynamics.

The second is motivated by the discussion of labor hoarding, and pro-cyclical productivity in the research on the relation between output and employment. Our benchmark model has assumed that labor productivity is exogenous; there is however substantial evidence is however that, perhaps due to labor hoarding, some of the movements in productivity are in fact endogenous. Thus, in contrast to our assumption, a positive realization of the noise shock may lead consumers to spend more, and lead in turn to an increase in productivity.

We consider these two extensions in turn.

5.1 Slow adjustment of consumption

To capture slow consumption adjustment, we adopt the simple specification, which replaces (5),

$$c_t = \beta c_{t-1} + (1-\beta) \lim_{j\to\infty} E_t[a_{t+j}].$$

In Table 4 we report the results from estimating this variant of the model, presented as a grid search over the value of the adjustment parameter β. The data seem to prefer a small but positive value of β.

β	ρ	σ_u	σ_ϵ	σ_ζ	σ_ν	ML
0	0.8785	0.0068	0.0008	0.0063	0.0086	-1073.3
0.1	0.87	0.0071	0.0009	0.0066	0.008	-1075.9
0.2	0.8591	0.0075	0.0011	0.007	0.0072	-1074.8
0.3	0.8412	0.0082	0.0013	0.0075	0.0062	-1068.8
0.4	0.7823	0.0092	0.002	0.0081	0.0035	-1057
0.5	0.6915	0.0107	0.0033	0.0089	0.0002	-1044.4
0.6	0.7126	0.013	0.0037	0.011	0.0003	-1018.2
0.7	0.6524	0.0177	0.0061	0.0143	0.0006	-976.7
0.8	0.6371	0.0272	0.0099	0.0217	0.0012	-910.9
0.9	0.648	0.0567	0.02	0.0456	0.0033	-796

Table 4: Maximum Likelihood Estimation: Slow Consumption Adjustment

5.2 Labor hoarding

[...]

5.3 Comparing true IRFs to estimated IRFs

[...]

6 Conclusions

- Methodologically: Limits of SVARs. Can do structural estimation. But then requires a model we truly believe. Estimation very non linear, and results depend a lot on specific assumptions.

- Empirically: Data quite consistent with smooth permanent shocks, and much short term action coming from transitory shocks and noise. Role of noise not well identified.
- Need to extend the model in many dimensions before having confidence in the conclusions. Distinguish between employment, output, consumption.
- If and when, can one interpret noise as animal spirits? How large is their contribution?
- Can in principle test for overoptimism, defined as a stronger response to expectations than implied by Kalman filtering. Can it realistically work? Not sure. (Hard to pinpoint in estimation the standard deviation of the noise, so it may be hard to estimate that additional coefficient as well.)

Appendix A. Relation of the model with the standard New Keynesian model

Consider a standard New Keynesian model, as laid out, e.g., in Gali (2008). Preferences are given by

$$E\sum_{t=0}^{\infty}\beta^{t}U\left(C_{t},N_{t}\right),$$

with

$$U\left(C_{t},N_{t}\right)=\log C_{t}-\frac{1}{1+\zeta}N_{t}^{1+\zeta},$$

where N_t are hours worked and C_t is a composite consumption good given by

$$C_{t}=\left(\int_{0}^{1}C_{j,t}^{\frac{\gamma-1}{\gamma}}dj\right)^{\frac{\gamma}{\gamma-1}},$$

$C_{j,t}$ is the consumption of good j in period t, and $\gamma>1$ is the elasticity of substitution among goods. Each good $j\in[0,1]$ is produced by a single monopolistic firm with access to the linear production function

$$Y_{j,t}=A_{t}N_{j,t}. \tag{11}$$

Productivity is given by $A_t=\exp a_t$ and a_t follows the process (1)-(3). Firms are allowed to reset prices only at random time intervals. Each period, a firm is allowed to reset its price with probability $1-\theta$ and must keep the price unchanged with probability θ. Firms hire labor on a competitive labor market at the wage W_t, which is fully flexible.

Consumers have access to a nominal one-period bond which trades at the price Q_t. The consumer's budget constraint is

$$Q_{t}B_{t+1}+\int_{0}^{1}P_{j,t}C_{j,t}dj=B_{t}+W_{t}N_{t}+\int_{0}^{1}\Pi_{j,t}dj, \tag{12}$$

where B_t are nominal bonds' holdings, $P_{j,t}$ is the price of good j, W_t is the nominal wage rate, and $\Pi_{j,t}$ are the profits of firm j. In equilibrium consumers choose consumption, hours worked, and bond holdings, so as to maximize their expected utility subject to (12) and a standard no-Ponzi-game condition. Nominal bonds are in zero net supply, so market clearing in the bonds market requires $B_t=0$. The central bank sets the short-term nominal interest rate, that is, the price of the one-period nominal bond, Q_t. Letting $i_t=-\log Q_t$, monetary policy follows the simple rule

$$i_{t}=i^{*}+\phi\pi_{t}, \tag{13}$$

where $i^* = -\log \beta$ and ϕ is a constant coefficient greater than 1.

Following standard steps, consumers' and firms' optimality conditions and market clearing can be log-linearized and transformed so as to obtain two stochastic difference equations which characterize the joint behavior of output and inflation in equilibrium. After substituting the policy rule we obtain:

$$\begin{aligned} y_t &= E_t\left[y_{t+1}\right] - \phi \pi_t + E_t\left[\pi_{t+1}\right], \\ \pi_t &= \kappa\left(y_t - a_t\right) + \beta E_t\left[\pi_{t+1}\right], \end{aligned}$$

where $\kappa \equiv (1+\zeta)(1-\theta)(1-\beta\theta)/\theta$ and where constant terms are omitted. As long as $\phi > 1$ this system has a unique locally stable solution where y_t and π_t are linear functions of the four exogenous state variables $a_t, x_{t|t}, x_{t-1|t}, z_{t|t}$,

$$\begin{pmatrix} y_t \\ \pi_t \end{pmatrix} = D_\kappa \begin{pmatrix} a_t \\ x_{t|t} \\ x_{t-1|t} \\ z_{t|t} \end{pmatrix}.$$

The matrix D_κ can be found using the method of undetermined coefficient as the solution to

$$\begin{bmatrix} 1 & \phi \\ -\kappa & 1 \end{bmatrix} D_\kappa = \begin{bmatrix} 0 & 0 & 0 & 0 \\ -\kappa & 0 & 0 & 0 \end{bmatrix} + \begin{bmatrix} 1 & 1 \\ 0 & \beta \end{bmatrix} D_\kappa \begin{bmatrix} 0 & 1+\rho & -\rho & \rho \\ 0 & 1+\rho & -\rho & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & \rho \end{bmatrix}.$$

The elements of D_κ are a continuous non-linear function of κ and some lengthy algebra (available on request) shows that

$$\lim_{\kappa \to 0} D_\kappa = \frac{1}{1-\rho} \begin{bmatrix} 0 & 1 & -\rho & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}.$$

Since $\kappa \to 0$ when $\theta \to 1$, this completes the argument.

Appendix B. Kalman filter

Let

$$C \equiv \begin{bmatrix} 1+\rho & -\rho & 0 \\ 1 & 0 & 0 \\ 0 & 0 & \rho \end{bmatrix}, D \equiv \begin{bmatrix} 1 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix},$$

and

$$\Sigma_1 \equiv \begin{bmatrix} \sigma_\epsilon^2 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & \sigma_\eta^2 \end{bmatrix}, \quad \Sigma_2 \equiv \begin{bmatrix} 0 & 0 \\ 0 & \sigma_\nu^2 \end{bmatrix}.$$

Then the process for $\xi_t \equiv (x_t, x_{t-1}, z_t)$ is described compactly as

$$\xi_t = C\xi_{t-1} + (\epsilon_t, 0, \eta_t)',$$

and the observation equation for the consumers is

$$(a_t, s_t) = D\xi_t + (0, \nu_t)'.$$

Let $P \equiv Var_{t-1}[\xi_t]$. The value of P is found solving the equation

$$P = C\left[P - PD'(DPD' + \Sigma_2)^{-1} DP\right] C' + \Sigma_1.$$

The matrixes A and B in the text are then given by:

$$\begin{aligned} A &= (I - BD)\,C, \\ B &= PD'(DPD' + \Sigma_2)^{-1}. \end{aligned}$$

Appendix C. Estimated consumption responses from an SVAR

Let u_t^c and u_t^a be the reduced form residuals which satisfy, by definition,

$$\begin{aligned} u_t^c &= c_t - E[c_t | a_{t-1}, c_{t-1}, ...], \\ u_t^a &= a_t - E[a_t | a_{t-1}, c_{t-1}, ...]. \end{aligned}$$

The identified temporary shock is given by

$$w_t = \beta_c u_t^c + \beta_a u_t^a,$$

where β_a and β_c are chosen so that

$$\lim_{j\to\infty} E\left[a_{t+j}|w_t, a_{t-1}, c_{t-1}, ...\right] - \lim_{j\to\infty} E\left[a_{t+j}|a_{t-1}, c_{t-1}, ...\right] = 0.$$

Given that w_t is a linear combination of $a_t, c_t, a_{t-1}, c_{t-1}, ...$ and all these variables are in $\mathcal{I}_t$ we can apply the law of iterated expectations to get

$$\lim_{j\to\infty} E\left[a_{t+j}|w_t, a_{t-1}, c_{t-1}, ...\right] = E[\lim_{j\to\infty} E\left[a_{t+j}|\mathcal{I}_t\right]|w_t, a_{t-1}, c_{t-1}, ...] = E\left[c_t|w_t, a_{t-1}, c_{t-1}, ...\right],$$

and

$$\lim_{j\to\infty} E\left[a_{t+j}|a_{t-1}, c_{t-1}, ...\right] = E[\lim_{j\to\infty} E\left[a_{t+j}|\mathcal{I}_t\right]|a_{t-1}, c_{t-1}, ...] = E\left[c_t|a_{t-1}, c_{t-1}, ...\right].$$

It follows that

$$E\left[c_t|w_t, a_{t-1}, c_{t-1}, ...\right] - E\left[c_t|a_{t-1}, c_{t-1}, ...\right] = 0,$$

which proves the proposition.

References

Beaudry, P., Portier, F., 2006, "Stock Prices, News, and Economic Fluctuations," *American Economic Review*, 96(4), 1293-1307.

Blanchard, O., Quah, D., 1989, The Dynamic Effects of Aggregate Demand and Aggregate Supply Disturbances, *American Economic Review*, 79(4), 654-673.

Galí, J., 1999, "Technology, Employment and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review*, 89(1), 249-271.

Galí, J., 2008, *Monetary Policy, Inflation and the Business Cycle: An Introduction to the New Keynesian Framework.* Princeton University Press.

Lorenzoni, G., 2008, "A Theory of Demand Shocks", mimeo MIT.

Quah, D., 1990, "Permanent and Transitory Movements in Labor Income: An Explanation for 'Excess Smoothness' in Consumption," *Journal of Political Economy*, 98(3), 449-75.

Rodriguez Mora, J., Schulstad, P., 2007, "The Effects of GNP announcements on Fluctuations of GNP Growth," *European Economic Review*, 51(8), 1922-1940.