

News, Noise, and Fluctuations: An Empirical Exploration

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

May 20, 2009

Abstract

We explore empirically models of aggregate fluctuations with two basic ingredients: agents form anticipations about the future based on noisy sources of information; these anticipations affect spending and output in the short run. Our objective is to separate fluctuations due to actual changes in fundamentals (news) from those due to temporary errors in the private sector's estimates of these fundamentals (noise). Using a simple model where the consumption random walk hypothesis holds exactly, we address some basic methodological issues and take a first pass at the data. First, we show that if the econometrician has no informational advantage over the agents in the model, structural VARs cannot be used to identify news and noise shocks. Next, we develop a structural Maximum Likelihood approach which allows us to identify the model's parameters and to evaluate the role of news and noise shocks. Applied to postwar U.S. data, this approach suggests that noise shocks play an important role in short-run fluctuations.

Keywords: Aggregate shocks, business cycles, vector autoregression, invertibility.

JEL Codes: E32, C32, D83

*IMF, MIT and NBER; MIT; MIT and NBER. We thank Daron Acemoglu, Marios Angeletos, Ricardo Caballero, Larry Christiano, Martin Eichenbaum, Matteo Iacoviello, Anna Mikusheva, Roberto Rigobon, Paul Schrimpf, Chris Sims, Mark Watson, Iván Werning and seminar participants at MIT, the NBER Summer Institute, the FED Board, University of Southern California, University of Maryland, EIEF, the AEA Meetings (San Francisco), Indiana University, the Richmond FED, and Brandeis University for very useful suggestions.

Introduction

A common view of the business cycle gives a central role to anticipations. Consumers and firms continuously receive information about the future, which sometimes is news, sometimes just noise. Based on this information, consumers and firms choose spending and, because of nominal rigidities, spending affects output in the short run. If ex post the information turns out to be news, the economy adjusts gradually to a new level of activity. If it turns out to be just noise, the economy returns to its initial state. Therefore, the dynamics of news and noise generate both short-run and long-run changes in aggregate activity. In this paper, we ask how aggregate time series can be used to shed light on this view of the business cycle.

We are interested in this view 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 coming here from a rational reaction to information about the future—, the role of demand in affecting output in the short run, together with the notion that in the long run output follows a natural path determined by fundamentals.

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 “transitory shocks”—by construction—lead to a transitory effect on activity in the short run. It is tempting to associate shocks with permanent effects to news and shocks with transitory effects to noise.

In this paper, we focus on a simple model which provides a useful laboratory to address two issues: a methodological one and a substantive one. First, can structural VARs indeed be used to recover news and noise shocks? Second, what is the role of news and noise shocks in short-run fluctuations?

On the first question, we reach a strong negative conclusion—one which came as an unhappy surprise for one of the coauthors. In models of expectation-driven fluctuations in which consumers solve a signal extraction problem, structural VARs can typically recover neither the 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.

To address the second question, we then turn to structural estimation, first using a

simple method of moments and then Maximum Likelihood. We find that our model fits the data well and gives a clear description of fluctuations as a result of three types of shocks: shocks with permanent effects on productivity, which build up slowly over time; shocks with temporary effects on productivity, which decay slowly; and shocks to consumers' signals about future productivity. All three shocks affect agents' expectations, and thus demand and output in the short run, and noise shocks are an important source of short-run volatility. In our baseline specification, noise shocks account for more than half of the forecast error variance at a yearly horizon, while permanent technology shocks account for less than one third. This result is somewhat surprising when compared with variance decompositions from structural VARs where transitory "demand shocks" often account for a smaller fraction of aggregate volatility at the same horizons and permanent technology shock capture a bigger share (e.g., Shapiro and Watson, 1989, and Galí, 1992). Our methodological analysis helps to explain the difference, showing why structural VARs may understate the contribution of noise/demand shocks to short-run volatility and overstate that of permanent productivity shocks.

Recent efforts to empirically estimate models of news-driven business cycles include Christiano, Ilut, Motto and Rostagno (2007) and Schmitt-Grohé and Uribe (2008). These papers follow the approach of Jaimovich and Rebelo (2006), modeling news as advanced, perfect information about shocks affecting future productivity. We share with those papers the emphasis on structural estimation. The main difference is that we model the private sector information as coming from a signal extraction problem and focus our attention on disentangling the separate effects of news and noise.

The problem with structural VARs emphasized in this paper is essentially an invertibility problem, also known as non-fundamentalness. There is a resurgence of interest in the methodological and practical implications of invertibility problems, see, e.g., Sims and Zha (2006) and Fernández-Villaverde, Rubio-Ramírez, Sargent and Watson (2007). Our paper shows that non-invertibility problem are endemic to models where the agents' uncertainty is represented as a signal extraction problem. This idea has also recently surfaced in models that try to identify the effects of fiscal policy when the private sector receives information on future policy changes (see Leeper, Walker and Yang, 2009).

The paper is organized as follows. Sections 1 and 2 present and solve the model. Section 3 looks at the use of structural VARs. Section 4 presents the results of our structural estimation. Section 5 explores a number of extensions and 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 postwar 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 have more information than just current and past productivity, we allow them to observe an additional signal about the permanent component of productivity. Having solved the signal extraction problem, and based on their expectations, agents 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 as the “permanent shock”, the “transitory shock”, and the “noise shock”. Permanent shock is a slight (and common) misnomer, as it refers to a shock whose effects build up gradually.

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 shocks with variances σ_ϵ^2 and σ_η^2 . Agents observe productivity, but not the two components separately.¹

For most of the paper, we assume that the univariate representation of a_t is a random walk

$$a_t = a_{t-1} + u_t, \quad (4)$$

with the variance of u_t equal to σ_u^2 , and restrict attention to 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, this assumption provides a surprisingly good starting point when looking at postwar U.S. data. As will be clear, however, none of our central results depends on this assumption.

In general, a given univariate process is consistent with an infinity of decompositions between a permanent and a transitory component with orthogonal innovations, as shown in Quah (1990, 1991). In our setup, there is a one-parameter family of processes (1)-(3) which deliver the univariate random walk (4). This is the family of processes that satisfy the following conditions:

$$\begin{aligned} \rho_x &= \rho_z = \rho, \\ \sigma_\epsilon^2 &= (1 - \rho)^2 \sigma_u^2, \quad \sigma_\eta^2 = \rho \sigma_u^2, \end{aligned}$$

for some $\rho \in [0, 1)$.²

Productivity may be the sum of a permanent process with small shocks that build up slowly and a transitory process with large shocks that decay slowly (high ρ , small σ_ϵ^2 and large σ_η^2), or it may be the sum of a permanent process which is itself close to a random walk and a transitory process close to white noise with small variance (low ρ , large σ_ϵ^2 and small σ_η^2). An econometrician who can only observe a_t cannot distinguish these cases. The sample variance of Δa_t gives an estimate of σ_u^2 , but the parameter ρ , and thus ρ_x , ρ_z , σ_ϵ^2 and σ_η^2 , are not identified. As we shall see, when consumers have some additional source of information

¹A similar process for technology, which combines level and growth rate shocks, has been recently used in an open economy context by Aguiar and Gopinath (2007). Boz, Daude and Durdu (2008) explore the role of different informational assumptions in that context.

²To prove this result, notice that, in general, (1)-(3) imply

$$Var[\Delta a_t] = \frac{1}{1 - \rho_x^2} \sigma_\epsilon^2 - \frac{2}{1 + \rho_z} \sigma_\eta^2,$$

and

$$Cov[\Delta a_t, \Delta a_{t-j}] = \rho_x^j \frac{1}{1 - \rho_x^2} \sigma_\epsilon^2 - \rho_z^{j-1} \frac{1 - \rho_z}{1 + \rho_z} \sigma_\eta^2 \text{ for all } j > 0.$$

Under the assumed parameter restrictions these yield $Var[\Delta a_t] = \sigma_u^2$ and $Cov[\Delta a_t, \Delta a_{t-j}] = 0$ for all $j > 0$.

on the permanent component x_t and the econometrician has access to consumption data, he will be able to identify ρ and the remaining parameters.

1.2 Consumption

We assume that consumption smoothing leads to the Euler equation

$$c_t = E[c_{t+1}|\mathcal{I}_t], \quad (5)$$

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, in which consumption is the only component of demand and output is fully determined by the demand side. 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

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

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.

Combining the last two equations gives

$$c_t = \lim_{j \rightarrow \infty} E_t[a_{t+j}]. \quad (6)$$

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, they receive a signal regarding the permanent component of the productivity process

$$s_t = x_t + \nu_t, \quad (7)$$

where ν_t is i.i.d. normal with variance σ_ν^2 . Moreover, consumers know the structure of the model, i.e., know ρ and the variances of the three shocks.

Finally, on the econometrician's side, we will consider both the case where the signal s_t is directly observable and the econometrician has access to time series for a_t , c_t and s_t , and

the case where only a_t and c_t are observed (as it will be the case in our empirical exercise). We will use \mathcal{I}_t^e to denote the econometrician's information set.

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_t , η_t , and ν_t . It is derived in two steps. First, we solve for consumption as a function of productivity expectations. From equations (1)-(3) and (6) above, we obtain

$$c_t = x_{t|t} + \frac{\rho}{1-\rho}(x_{t|t} - x_{t-1|t}). \quad (8)$$

Recall that $x_{t|t}$ and $x_{t-1|t}$ denote the consumers' expectations about the unobservable states x_t and x_{t-1} .

Second, we derive the dynamics of the expectations in (8) using the Kalman filter. 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 update their beliefs applying the Kalman filter:

$$\begin{bmatrix} x_{t|t} \\ x_{t-1|t} \\ z_{t|t} \end{bmatrix} = A \begin{bmatrix} x_{t-1|t-1} \\ x_{t-2|t-1} \\ z_{t-1|t-1} \end{bmatrix} + B \begin{bmatrix} a_t \\ s_t \end{bmatrix} \quad (9)$$

where the matrices A and B depend on the underlying parameters (see Appendix B).

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

Figure 1 shows the impulse responses of consumption and productivity computed using parameters in line with the estimates obtained later, in Section 4. The time unit is the quarter. The parameter ρ is set to 0.89, implying slowly building permanent shocks and slowly decaying transitory shocks. The standard deviation of productivity growth, σ_u , is set to 0.67%. These values for ρ and σ_u yield standard deviations of the two technology shocks, σ_ϵ and σ_η , equal to 0.07% and 0.63%, respectively. The standard deviation of the noise shock, σ_ν , is set to 0.89%, implying a fairly noisy signal.

In response to a one standard deviation increase in ϵ_t , a permanent technology shock, productivity builds up slowly over time—the implication of a high value for ρ . Consumption




Figure 1: Impulse Responses to the Three Shocks

also increases slowly. This reflects the fact that the standard deviations of the transitory shock η_t and of the noise shock ν_t are both large relative to the standard deviation of ϵ_t . Thus, it takes a long time for consumers 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_t , 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_t , 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, the response turns briefly negative, before returning to normal. By assumption, productivity does not change, so employment initially increases, to return to normal over time.

2.1 Innovations representation

Our assumptions make it easy to derive the *innovations representation* of the processes for consumption and productivity.³ In particular, rearranging (8), we obtain

$$(1 - \rho)c_t = x_{t|t} - \rho x_{t-1|t}. \quad (10)$$

Writing the corresponding expression for c_{t-1} and taking differences side by side, we obtain

$$c_t = c_{t-1} + u_t^c, \quad (11)$$

with

$$u_t^c = \frac{1}{1 - \rho}(x_{t|t} - x_{t-1|t-1}) - \frac{\rho}{1 - \rho}(x_{t-1|t} - x_{t-2|t-1}).$$

³See Anderson, Hansen, McGrattan, and Sargent (1996) for general results on the existence of an innovations representation.

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

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

Using (10), lagged one period, we then obtain

$$a_t = \rho a_{t-1} + (1 - \rho) c_{t-1} + u_t^a, \quad (12)$$

with

$$u_t^a = x_t - x_{t-1|t-1} - \rho(x_{t-1} - x_{t-2|t-1}) + \eta_t.$$

To show that u_t^c and u_t^a in (11) and (12) are indeed innovations take expectations and use (2) to obtain

$$\begin{aligned} E_{t-1}[u_t^c] &= \frac{1}{1 - \rho} E_{t-1}[\epsilon_t] = 0, \\ E_{t-1}[u_t^a] &= E_{t-1}[\epsilon_t + \eta_t] = 0. \end{aligned}$$

This shows that u_t^c and u_t^a are innovations with respect to the consumers' information. Turning to the econometrician, we can assume that the econometrician observes (c_t, a_t, s_t) or just (c_t, a_t) . In either case the econometrician has (weakly) less information than the consumer and the law of iterated expectations implies $E[u_t^c | \mathcal{I}_t^e] = 0$ and $E[u_t^a | \mathcal{I}_t^e] = 0$. Therefore, u_t^c and u_t^a represent innovations for consumption and productivity both in a reduced form VAR in (c_t, a_t, s_t) and in a reduced form VAR in (c_t, a_t) .

Note that, under our assumptions, the univariate representations of both productivity and consumption are random walks. For productivity this follows from our assumptions on the productivity process, for consumption it follows from the behavioral assumption (5). When we move to multivariate representations including c_t and a_t , past productivity does not help predict consumption, but, as (12) shows, past consumption typically helps to predict productivity as it captures the consumers' information on the permanent component x_t .⁴

⁴The special case in which consumption does not help to predict productivity is $\rho = 0$. As we shall see below, in this case a_t and c_t are perfectly collinear, so, given a_{t-1} , c_{t-1} provides no extra information on a_t . In this case, the innovations representation is not unique, as (12) can be replaced, for example, by $a_t = a_{t-1} + u_t^a$.

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.

The answer to this question is, generally, no. The basic intuition is the following: if consumption is a random walk given the consumers' information sets, then an econometrician with access to the same information, or less, cannot identify any shock that has a transitory effect on consumption based on the reduced form VAR innovations at time t . If the econometrician could, so would the agents. But then they would optimally choose a consumption path that does not respond to these identified shocks.

In the rest of this section we flesh out this intuition and show how it leads to a non-invertibility problem. We begin from two special cases, the case where the signal s_t is perfectly informative, $\sigma_\nu = 0$, and the case where it is completely uninformative, $\sigma_\nu = \infty$. In both cases, noise shocks do not affect the consumption and productivity dynamics, so we can focus on the econometrician's problem of recovering the two shocks ϵ_t and η_t from the bivariate time series (c_t, a_t) .

3.1 A perfectly informative signal

If the signal is perfectly informative, 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 (11) and (12) 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 does work. Imposing the long-run restriction that only one of the shocks has a permanent effect on consumption and productivity, we can recover ϵ_t and η_t , and their dynamic effects.

3.2 An uninformative signal

If, instead, the signal is uninformative, the consumers rely only on current and past productivity to forecast future productivity. 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 . That is, 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 importance of permanent and transitory productivity shocks, 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_t will show a one-time permanent increase; the impulse responses of productivity and consumption to η_t will be identically equal to zero.

However, in this case the decomposition between temporary and permanent shocks is essentially irrelevant, given that no information is available to ever separate the two. We might as well take the random walk representation of productivity as our primitive productivity process and just interpret u_t as the single, permanent shock. With this interpretation, one can safely adopt a structural VAR approach.

3.3 The general case

In the two special cases just considered, a structural VAR approach seems to work, albeit for very different reasons: In the first, we can exploit the perfect information of the consumers to separate permanent and transitory shocks. In the second, we can ignore the “true” productivity process and just focus on the observable random walk for productivity.

Unfortunately, once we move away from these special cases and have a partially informative signal, a structural VAR approach fails. In the general case, unlike in the first case, the consumers’ information at time t is not sufficient to exactly recover the shocks. At the same time, unlike in the second case, consumption reflects some information on the transitory and permanent components of productivity, so we cannot ignore their underlying dynamics.

Now the model features three shocks, ϵ_t , η_t and ν_t , so we consider the econometrician’s problem of recovering these three shocks from the trivariate time series (c_t, a_t, s_t) . The econometrician runs a reduced form VAR in (c_t, a_t, s_t) and obtains the reduced form innovations (u_t^c, u_t^a, u_t^s) . He then tries to use some identification restriction to map the reduced form innovations into the economic shocks. An identified shock will correspond to a linear com-

bination of reduced form innovations. The next proposition characterizes the shape of the estimated responses of consumption to any identified shock.

Proposition 1 *Suppose that the econometrician observes (c_t, a_t, s_t) . Then, the estimated impulse response of c_t to any identified shock from a structural VAR will be, asymptotically, either permanent and flat or zero.*

Comparing this result with the impulse responses obtained in Figure 1 immediately shows that a structural VAR will be, in general, unable to recover the model’s responses to our three shocks, given that none of them leads to a flat consumption response.

Why does the structural VAR fail? Suppose there was an identified structural shock that could be mapped into the noise shock of the model. That means that there would be a linear combination of reduced form innovations at time t that can be used to forecast the transitory increase in consumption in panel (c) of Figure 1. The consumers have access to all the data used by the econometrician to construct the innovations at time t : they know the model’s parameters and they have observed all variable realizations up to time t . Therefore, they must also be able to forecast this transitory fluctuation in consumption. But this would violate consumption smoothing. Therefore, the consumption response to any identified shock must be flat.

This is not a problem in the special case where consumers have a perfectly informative signal, because in that case the impulse responses in the model coincide with the ones in Proposition 1: permanent and flat response to ϵ_t and zero response to η_t . The same is true in the special case of an uninformative signal, if we limit ourselves to recovering responses to the shock u_t . In the general case, however, the impulse responses are richer than those in Proposition 1. Moreover, as we shall see in Section 4, the data contain enough information to estimate these responses. The problem is that a structural VAR approach tries to get there by exactly recovering the shocks at time t from the observables up to that period, and this is not feasible in the general case.

Notice that our specific assumptions on the productivity process and on the informational structure are not crucial for Proposition 1. In fact, the result can be extended to any process for a_t and any signal process, as long as the consumption process is well defined and satisfies $c_t = \lim_{j \rightarrow \infty} E[a_{t+j} | \mathcal{I}_t]$.

One could enrich the model, e.g., adding preference shocks and allowing for changes in the real interest rate, so as to relax the random walk hypothesis for consumption. However, the essence of the argument remains: noise shocks that lead to transient “mistakes” by consumers

cannot be detected using information available to consumers at date t . A structural VAR identification scheme can only use that information and is bound to fail.

Proposition 1 clearly extends to the case where the econometrician only observes the bivariate series (c_t, a_t) . Given that this will be the information set used in our empirical exercise in Section 4, it is useful to analyze this case in more detail. In particular, we can use a numerical example to further investigate the direction of the bias in the estimated impulse responses.

Figure 2 shows the estimated impulse responses to the shocks with permanent and transitory effects obtained from structural VAR estimation, together with the true impulse responses to the three underlying shocks. The underlying parameters are the same as for Figure 1. The estimated impulse responses are obtained by generating a 10,000-period time series for consumption and productivity using the true model and running a structural VAR on it. The structural VAR is identified by imposing a long-run restriction which distinguishes two orthogonal shocks: one with permanent effects on output and one with only transitory effects.

Look first at the true and estimated responses of productivity to a shock with permanent effects. The solid line in the top left quadrant plots the true response to a permanent technology shock, which replicates that in Figure 1, namely a small initial effect, followed by a steady buildup over time. The dashed line gives the estimated response from the structural VAR 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 response shows a one-time increase, with no further build up over time).

Turn to the true and estimated responses of consumption to a permanent shock in the bottom left quadrant. The solid line again replicates the corresponding response in Figure 1, showing a slow build-up of consumption over time. The dashed line shows the estimated response, namely a one-time response of consumption with no further build up over time.

The right quadrants show the true and estimated responses to shocks with transitory effects on output. The solid lines show the true responses to a transitory technology shock (thick line) and to a noise shock (thin line). The dashed lines give the estimated response to the single transitory shock from the structural VAR. They show that the estimated response of productivity to a transitory shock is close to the true response to a transitory technology shock, but the estimated response of consumption is equal to zero.

In short, the responses from the structural VAR overstate the initial response of productivity and consumption to permanent shocks, and thus give too much weight to these shocks




Figure 2: True and SVAR-based estimated impulse responses

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.

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

The result above suggests two potential ways out, both based on the possibility that the econometrician may have access to more information than the agents, either at time t or later.

First, 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, it suggests that the bias from SVAR estimation may be reduced when using revised series rather than originally published series.⁵ The dispersed information model in Lorenzoni (2009) goes in this direction, by assuming that consumers do not have access to real time information on aggregate output, but only to noisy local information. Under that assumption it is possible to map the noise shock in that model to the transitory shock from an identified VAR. However, also in models with dispersed information, once we enrich the consumers' information set, the problem raised here is bound to reappear.

The need for superior information on the econometrician's side, suggests a second way out. In the end, the econometrician always has access to superior information, as he can observe future realizations of variables that the consumer did not observe at time t . Then one may hope that a combination of past and future data may be used to identify current shocks. More formally, the traditional invertibility problem is that the map from the economic shocks to the shocks in the VAR may not have an inverse that is one-sided in nonnegative powers of the lag operator. Maybe adding a sufficient number of lead terms an inverse can be found? Unfortunately, the answer is no. As we will show numerically in Section 4.5, even having

⁵A 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.

access to an infinite sequence of past and future data the econometrician is never able to exactly recover the values of the shocks.

3.5 What does the structural VAR deliver?

A different way of looking at the problem is to understand what is the correct interpretation of the identified shocks that the structural VAR delivers. It turns out that the structural VAR allows us to recover the process for a_t in its innovations representation. Namely, the process for a_t can be equivalently represented by the state-space system:

$$\hat{x}_t = \hat{x}_{t-1} + v_t^1 \tag{13}$$

$$a_t = \rho a_{t-1} + (1 - \rho)\hat{x}_{t-1} + v_t^2. \tag{14}$$

To prove the equivalence it is sufficient to define $\hat{x}_t \equiv c_t$, and use the results in Section 2.1, substituting v_t^1 for u_t^c and v_t^2 for u_t^a .

But then why not start directly from (13)-(14) as our model for productivity dynamics and give consumers full information on the state \hat{x}_t ? One reason why (13)-(14) is not particularly appealing as a primitive model is that the disturbances v_t^1 and v_t^2 in the innovation representation above are not mutually independent, and thus are hard to interpret as primitive shocks. In particular, our signal extraction model implies that v_t^1 and v_t^2 are positively correlated and their correlation is higher the higher the underlying value of σ_ν . As we shall see in the next section, this positive correlation is indeed observed in the data. Our informational assumptions provide a rationale for it.

Going back to structural VARs, a long-run identifying restriction will lead us to identify v_t^1 as the permanent technology shock and will give a linear combination of v_t^1 and v_t^2 as the temporary shock. For some purposes, this representation may be all we are interested in. Clearly, that is not the case if we are trying to analyze the role of noise shocks in fluctuations.

4 Structural estimation

We now turn to structural estimation, proceeding in two steps. For our benchmark model structural estimation is particularly easy, and all parameters can be obtained matching a few moments of the model to the data; thus we start with it. For more general processes however, one must use maximum likelihood. We show how it can be done, show estimation results for our benchmark model and compare them to those obtained by matching moments.

4.1 Matching moments

In general, structural estimation allows us to exploit the cross-equation restrictions implied by the model to achieve identification. Equation (12), our reduced form equation for productivity, provides a good example of this principle: estimating this equation by OLS, allows us immediately to recover the parameter ρ . Moreover, σ_u^2 can be estimated by the sample variance of Δa_t . Having estimates for ρ and σ_u^2 , we immediately get estimates for σ_ϵ^2 and σ_η^2 .

Although identification is particularly simple here, the point holds more generally. In the class of models considered here, identification can be achieved exploiting two crucial assumptions: some behavioral assumption which links consumption (or some other endogenous variables) to the agents' expectations about the future, here equation (6), and an assumption of rational expectations.⁶

How well does our reduced form benchmark model (11)-(12) fits 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 provides a good starting point.

Throughout this section, we only use time series for a_t and c_t . We construct the productivity variable as the logarithm of the ratio of GDP to employment and 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, productivity and consumption have different growth rates over the sample (0.34% per quarter for productivity, versus 0.46% 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, in what follows, we allow for a secular drift in the consumption-to-productivity ratio (equal to 0.46%-0.34%) and remove it from the consumption series.⁷

The basic characteristics of the time series for productivity and consumption are presented in Table 1. Lines 1 and 2 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

⁶The use of behavioral assumptions as identification assumptions to estimate an underlying exogenous process, connects our paper to a large body of work on household income dynamics. See, for example, Blundell and Preston (1998), who use the permanent income hypothesis as an identification assumption to decompose the household income process into transitory and permanent components.

⁷We 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. The reason why we concentrate on the sample 1970:1 to 2008:1 is precisely because, with longer samples, we are less confident that this approach does a satisfactory job at accounting for low frequency changes in the consumption-to-productivity ratio. When we turn to the variance decomposition, we will show that our results are robust to extending the sample.

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 consumption, equation (12), which can be rewritten as the cointegrating regression:

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

Line 3 shows the results of estimating 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 significant coefficients on the lagged rates of change of both variables.

| Line | Dependent variable: | $\Delta a(-1)$ | $\Delta c(-1)$ | $(c - a)(-1)$ |
|------|---------------------|----------------|----------------|---------------|
| 1 | Δa | -0.06 (0.09) | | |
| 2 | Δc | | 0.24 (0.08) | |
| 3 | Δa | | | 0.05 (0.03) |
| 4 | Δa | -0.21 (0.10) | 0.32 (0.12) | 0.03 (0.02) |
| 5 | $\Delta(8)a$ | | | 0.03 (0.15) |
| 6 | $\Delta(20)a$ | | | 0.31 (0.30) |
| 7 | $\Delta(40)a$ | | | 0.98 (0.43) |

Table 1: Consumption and Productivity Regressions.

Note: Sample: 1970:1 to 2008:1. $\Delta(j)a \equiv a(+j - 1) - a(-1)$. In parenthesis: robust standard errors computed using the Newey-West window and 10 lags.

Our model's dynamic implications on the relation between consumption and productivity can be extended to longer horizons. Specifically, (12) can be extended to obtain the following cointegrating regression, which holds for all $j \geq 0$,⁸

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

⁸This is obtained by induction. Suppose it is true for j , that is, $E_t[a_{t+j}] = (1 - \rho^j)c_t + \rho^j a_t$. Taking expectations at time $t - 1$ on both sides yields

$$\begin{aligned} E_{t-1}[a_{t+j}] &= (1 - \rho^j) E_{t-1}[c_t] + \rho^j E_{t-1}[a_t] \\ &= (1 - \rho^j) c_{t-1} + \rho^j ((1 - \rho) c_{t-1} + \rho a_{t-1}) \\ &= (1 - \rho^{j+1}) c_{t-1} + \rho^{j+1} a_{t-1}, \end{aligned}$$

the second equality follows from (5) and (12), the third from rearranging.

where $u_t^{a,j}$ is a disturbance uncorrelated to the econometrician's information at date t . Thus, according to the model, a larger consumption-productivity ratio should forecast higher future productivity growth at all horizons and the coefficient in this regression should increase with the horizon. Lines 5 to 7 explore this implication. We correct for the presence of autocorrelation due to overlapping intervals by using Newey-West standard errors. These results are roughly consistent with the model predictions, and all point to relatively high values for ρ : the point estimates implicit in lines 3, 5, 6 and 7 are, respectively, 0.95, 0.996, 0.98 and 0.91. The maximum likelihood approach below will use all the model restrictions to produce a single estimate of ρ , for now we just take the estimate from line 3, $\rho = 0.95$.

The standard deviation σ_u can be estimated directly from the univariate representation of a_t as the sample mean squared deviation of Δa_t , giving a point estimate $\sigma_u = 0.67\%$. Together with $\rho = 0.95$, this implies $\sigma_\epsilon = 0.03\%$ and $\sigma_\eta = 0.65\%$. 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.

Recovering the variance of the noise shock is less straightforward, but it can be done matching another moment: the correlation coefficient between the reduced form innovations u_t^c and u_t^a . In particular, numerical results show that, given the remaining parameters, this moment is an increasing function of σ_ν . Therefore, we recover this parameter by matching the correlation in the data. The coefficient of correlation between Δc and the residual of the regression on line 3 (corresponding, respectively to u_t^c and u_t^a) is equal to 0.52. If the signal was perfectly informative this correlation would be equal to 0.05, while if the signal had infinite variance it would be 1.⁹ Therefore, the observed correlation is consistent with the model and yields a fairly large standard deviation of the noise shock, $\sigma_\nu = 2.1\%$.

The fact that we are able in our benchmark model to recover all the model parameters by matching a few moments from the data, is clearly a special case. It is thus useful to develop a general approach, which can be applied to any specification of productivity or consumption behavior. We now discuss this approach, and then return to the data.

4.2 Maximum Likelihood

To estimate a model where consumers face a non trivial signal extraction problem, one can, generally, proceed in two steps.¹⁰

⁹These bounds can be derived from the analysis in Sections 3.1 and 3.2. To obtain the first, some algebra shows that under full information $Cov[u_t^c, u_t^a]/\sqrt{Var[u_t^c]Var[u_t^a]} = (1-\rho)/\sqrt{(1-\rho)^2 + \rho}$. The second bound is immediate.

¹⁰More detailed derivations are provided in Appendices B and D.

- 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 (7), and Kalman filtering gives us the updating equation (9).
- Next, take the point of view of the econometrician, 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 (9) becomes part of the description of the state's dynamics. The observation equations for the econometrician are now (1) and (10), where the second links consumption (observed by the econometrician), to consumers' expectations. The econometrician's Kalman filter is then 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 to 0.99.¹¹ For each value of ρ , we find the values of the remaining parameters that maximize the likelihood function and in the last column we report the corresponding likelihood value. The table shows that the likelihood function has a well-behaved maximum at $\rho = 0.89$, on line 6. The corresponding values of σ_ϵ and σ_η are 0.07% and 0.63%, respectively. The standard deviation of the noise shock σ_ν is 0.89%.

Relative to the moment matching approach in Section 4.1, the Maximum Likelihood approach uses all the implicit restrictions imposed by the model on the data generating process. This explains the difference between the estimates on line 6 of Table 2 and those obtained in Section 4.1. In particular, the Maximum Likelihood approach favors smaller values of ρ and σ_ν . However, if we look at line 8 of Table 2, we see parameters much closer to those in Section 4.1 and the likelihood gain from line 8 to line 6 is not too large. In other words, the data are consistent with a range of different combinations of ρ and σ_ν . When we look at the model's implications in terms of variance decomposition, we will consider different values in this range.

A 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

¹¹For all our Maximum Likelihood estimates we used Dynare (v.3), which allows for the use of matrices in the model section of the code. Our observables are first differences of labor productivity and consumption, so we use a diffuse Kalman Filter to initialize the variance covariance matrix of the estimator (a variance-covariance matrix with a diagonal of 10).

| Line | ρ | σ_u | σ_ϵ | σ_η | σ_ν | ML |
|------|--------|------------|-------------------|---------------|--------------|----------------|
| 1 | 0.00 | 0.0067 | 0.0067 | 0.0000 | 0.0089 | $-3 * 10^{12}$ |
| 2 | 0.25 | 0.0183 | 0.0137 | 0.0092 | 0.0000 | 859.2 |
| 3 | 0.50 | 0.0102 | 0.0051 | 0.0072 | 0.0000 | 980.5 |
| 4 | 0.70 | 0.0077 | 0.0023 | 0.0065 | 0.0026 | 1042.6 |
| 5 | 0.80 | 0.0071 | 0.0014 | 0.0064 | 0.0056 | 1064.5 |
| 6 | 0.89 | 0.0067 | 0.0007 | 0.0063 | 0.0089 | 1073.2 |
| 7 | 0.90 | 0.0067 | 0.0007 | 0.0064 | 0.0099 | 1073.1 |
| 8 | 0.95 | 0.0068 | 0.0003 | 0.0066 | 0.0234 | 1072.2 |
| 9 | 0.99 | 0.0063 | 0.0001 | 0.0063 | 0.0753 | 1068.5 |

Table 2: Maximum Likelihood Estimation: Benchmark Model

freely estimated. The estimation results are reported in Table 3 and are quite close to those obtained under the random walk assumption.

| | Estimate | Standard error |
|-------------------|----------|----------------|
| ρ_x | 0.8879 | 0.0478 |
| ρ_z | 0.8878 | 0.0474 |
| σ_η | 0.0065 | 0.0004 |
| σ_ϵ | 0.0007 | 0.0003 |
| σ_ν | 0.0090 | 0.0052 |
| ML | 1073.3 | |

Table 3: Maximum Likelihood Estimation: Unconstrained Model

4.3 Variance decomposition

What do our results imply in terms of the dynamic effects of the shocks and of variance decomposition? If we use the estimated parameters from the benchmark model (line 6 in Table 2), the dynamic effects of each shock are given in Figure 1 and were discussed in Section 2: 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.

Figure 3 presents the variance decomposition, plotting the contribution of the three shocks to forecast error variance, from 1 to 20 quarters ahead. The main result is that noise shocks are an important source of short run volatility, accounting for more than 70% of consumption volatility at a 1-quarter horizon and more than 50% at a 4-quarter horizon,




Figure 3: Variance Decomposition: Benchmark Model

while permanent technology shocks play a smaller role, having almost no effect on quarterly volatility and explaining less than 30% at a 4-quarter horizon. It is interesting to compare this result to traditional SVAR exercises, such as Shapiro and Watson (1989) and Gali (1992), where demand shocks typically explain a smaller fraction of aggregate volatility and permanent technology shocks play a bigger role. The analysis in Section 3 helps to explain these differences, by showing that, asymptotically, a SVAR is biased towards assigning 100% of the volatility to the permanent shock.

In Table 4, we report the results of some robustness checks. On each line, we report the fraction of consumption variance due to the noise shock at a 1, 4 and 8-quarter horizon, for different parameter values. Line 1 corresponds to our benchmark estimation. Line 2 reports the results obtained by setting ρ at a higher level and choosing the remaining parameters by maximum likelihood (line 8 of Table 2). The variance decomposition at short horizons is not very different, but noise shocks turn out to be more persistent under this parametrization and explain a much bigger fraction of variance at a 8-quarter horizon. On line 3 we report

the parameters obtained when estimating our model on a longer sample, 1948:1 to 2008:1. With this data set the estimate of ρ is larger and we obtain results analogous to the ones on line 2.

Finally, in lines 4 and 5 we experiment with changing only the volatility of noise shocks, keeping the other parameters fixed. In particular, relative to the benchmark, we first decrease and then increase σ_ν by one standard deviation (which is 0.0034 in our maximum likelihood estimate). Interestingly, it is the lower value of σ_ν that leads to the largest amount of noise-driven volatility. A lower σ_ν makes the signal s_t more precise, so consumers rely on it more. In our range of parameters, this leads to greater short-run volatility.

| Line | | Parameters | | | Noise-driven variance (fraction) | | |
|------|----------------------|------------|------------|--------------|----------------------------------|-----------|-----------|
| | | ρ | σ_u | σ_ν | 1 Quarter | 4 Quarter | 8 Quarter |
| 1 | benchmark | 0.89 | 0.0067 | 0.0089 | 0.75 | 0.53 | 0.23 |
| 2 | high ρ | 0.95 | 0.0068 | 0.0234 | 0.71 | 0.68 | 0.58 |
| 3 | sample 1948:1-2008:1 | 0.96 | 0.0099 | 0.0382 | 0.73 | 0.71 | 0.64 |
| 5 | low σ_ν | 0.89 | 0.0067 | 0.0055 | 0.82 | 0.46 | 0.17 |
| 4 | high σ_ν | 0.89 | 0.0067 | 0.0123 | 0.68 | 0.53 | 0.26 |

Table 4: Variance Decomposition: Robustness Checks

4.4 Recovering the states: retrospective history

So far we have focused on using structural estimation to estimate the model's parameters. Now we turn to the question: what information on the unobservable states and on the shocks can be recovered from structural estimation? We begin with the states.

Using the Kalman smoother it is possible to form Bayesian estimates of the state vector ξ_t^E using the full time series available and obtain a retrospective history of the U.S. business cycle. The top panel of Figure 4 plots estimates for the permanent component of productivity x_t obtained from our benchmark model. The solid line correspond to x_t , the dashed line to the consumers' real time estimate of the same variable $x_{t|t}$. Notice that both x_t and $x_{t|t}$ are unobservable states for the econometrician, so the two lines correspond to the Bayesian estimates of the respective state (see Appendix D).

Looking first at medium-run movements, the model identifies a gradual adjustment of consumers' expectations to the productivity slowdown in the 70s and a symmetric gradual adjustment in the opposite direction during the faster productivity growth after the mid 90s. Around these medium-run trends, temporary fluctuations in consumers' expectations




Figure 4: Smoothed estimates of the permanent component of productivity, of long-run productivity, and of consumers' real time expectations

Top panel: smoothed estimate of x_t (solid line) and of $x_{t|t}$ (dashed line)

Bottom panel: smoothed estimate of $x_{t+\infty}$ (solid line) and of $x_{t+\infty|t}$ (dashed line)

produce short-run volatility.

To gauge the short-run effects of expectational errors, however, the consumers' expectations of x_t are not sufficient, given that consumers project future growth based on their expectations of both x_t and x_{t-1} . For this reason, in the bottom panel of Figure 4, we plot the smoothed series for the consumers' real time expectations regarding long-run productivity, $x_{t+\infty|t} = (x_{t|t} - \rho x_{t-1|t}) / (1 - \rho)$, and compare it to the smoothed series for $x_{t+\infty}$. The model generates large short-run consumption volatility out of temporary changes in consumers' expectations of future productivity. Sometimes these changes occur when consumers' overstate current x_t (e.g., at the end of the 80s), other times when consumers slowly catch up to an underlying productivity acceleration and understate x_{t-1} (e.g., at the end of the 90s). Obviously, the model is too stylized to give a credible account of all cyclical episodes. For example, given the absence of monetary policy shocks the recession of 1981-82 is fully attributed to animal spirits.

The Kalman smoother also tells us how much information on the unobservable states is




Figure 5: RMSE of the estimated states at time t using data up to $t + j$

contained in past and future data. In particular, in Figure 5 we plot the root mean squared errors (RMSE) of the smoothed estimates of x_t and z_t , when data up to $t + j$ are available, for $j = 0, 1, 2, \dots$. Formally, these RMSE correspond to the square root of $E_{t+j}[(x_t - E_{t+j}[x_t])^2]$, and can be computed using two different information sets: the econometrician's, which only includes observations of c_t and a_t , and the consumer's, which also includes s_t . For simplicity, we compute RMSE at the steady state of the Kalman filter, that is, assuming the forecaster has access to an infinite series of data, from $-\infty$ to $t + j$. In this case, the econometrician's information set coincides with the consumer's, that is, the econometrician can back up the current value of s_t perfectly from current and past observations of c_t and a_t . Although we have not established this result analytically, it holds numerically in all our examples: the computed RMSE of the econometrician's estimate of s_t goes immediately to zero at $j = 0$. This implies that, in our model, with a sufficiently long data set, the direct observation of s_t does not add much to the econometrician's ability to recover the unobservable states (or the shocks).

Figure 5 shows that the contemporaneous estimate of the current state x_t has a standard deviation of 0.44%. By using future data, this standard deviation almost halves, to 0.28%. However, most of the relevant information arrives in the first six quarters, after that, there are minimal gains in the precision of the estimate.




Figure 6: Smoothed estimates of the shocks

4.5 Recovering the shocks: more on invertibility

Turning to the shocks, we know from our discussion of structural VARs 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. However, this does not mean that the data contain no information on the shocks. In particular, using the Kalman smoother the econometrician can form Bayesian estimates on ϵ_t , η_t , and ν_t using the entire time series available. Figure 6 plots these estimates for our benchmark model. As for the states, in Figure 7 we report the RMSE of the estimated shocks as a function of the number of leads available. To help the interpretation, each RMSE is normalized dividing it by the ex ante standard deviation of the respective shock (σ_ϵ , σ_η , and σ_ν).

Notice that if the model was invertible, the RMSE would be zero at $j = 0$. The fact that all RMSE remain bounded from zero at all horizons shows that even an infinite data set would not allow us to recover the shocks exactly.

The transitory shock η_t is estimated with considerable precision already on impact and the precision of its estimate almost doubles in the long run. The noise shock ν_t is less precisely estimated, but the data still tell us a lot about it, giving us an RMSE which is




Figure 7: Normalized RMSE of the estimated shocks at time t using data up to $t + j$

about 1/3 of the prior uncertainty in the long run. The shock that is least precisely estimated is the permanent shock ϵ_t . Even with an infinite series of future data, the residual variance is about 94% of the prior uncertainty on the shock.

How do we reconcile the imprecision of the estimate of ϵ_t with the fact that we have relatively precise estimates of the state x_t , as seen in Figure 5? The explanation is that the econometrician can estimate the cumulated effect of permanent productivity changes by looking at productivity growth over longer horizons, but cannot pinpoint the precise quarter in which the change occurred. Therefore, it is possible to have imprecise estimates of past ϵ_t 's, while having a relatively precise estimate of their cumulated effect on x_t . This also helps to explain the high degree of autocorrelation of the estimated permanent shocks in Figure 6. The smoothed estimates of ϵ_t in consecutive quarters tend to be highly correlated, as the econometrician does not know to which quarter to attribute an observed permanent change in productivity. Notice that the autocorrelation of the estimated shocks is not a rejection of the assumption of i.i.d. shocks, but purely a reflection of the econometrician's information. In fact, performing the same estimation exercise on simulated data delivers a similar degree of autocorrelation as the one obtained from actual data.

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 productivity process, the information structure, and the behavior of consumers. To explore how sensitive are the estimated parameters to the specific assumptions, we consider two extensions.

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 follows an exact 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, on line 2 of Table 1, shows evidence of richer dynamics.

Here we try two approaches. The first is to allow for some time variation in the real interest rate by turning to a standard New Keynesian model with Calvo pricing. Such a model is described in Appendix A and leads to a process for consumption (and output) of

the form

$$c_t = d_1 a_t + d_2 x_{t|t} + d_3 x_{t|t-1} + d_4 z_{t|t} \quad (15)$$

where the coefficients d are non-linear functions of the following parameters: the discount factor β , a parameter ϕ , reflecting the response of the nominal interest rate to inflation in the monetary policy rule, and a parameter κ , capturing the degree of nominal and real rigidities in price setting. We set β at 0.99 and estimate the remaining parameters by Maximum Likelihood, following the same steps laid out in 4.2. The results are reported in Table 5.

| | Estimate | Standard error |
|--------------|----------|----------------|
| κ | 0.0011 | 0.0004 |
| ϕ | 1.4436 | 0.1403 |
| ρ | 0.8780 | 0.0225 |
| σ_u | 0.0067 | 0.0004 |
| σ_ν | 0.0065 | 0.0019 |
| ML | 1073.8 | |

Table 5: Maximum Likelihood Estimation: standard New Keynesian model

Notice that the data prefer a very low value for κ , so the implications of the New Keynesian model are very close to those of the benchmark model. In particular, the implied values of the coefficients in (15) are

$$d_1 = 0.0016, \quad d_2 = 7.9250, \quad d_3 = -6.9266, \quad d_4 = 0.0359,$$

while, in our benchmark model, given the same $\rho = 0.878$, the corresponding values would be $d_2 = 1/(1-\rho) = 8.1967$, $d_3 = -\rho/(1-\rho) = -7.1967$ and zeros for d_1 and d_4 . The implied impulse responses are thus close to the ones in Section 2.

To capture slow consumption adjustment, we then try an alternative specification of consumption behavior, which incorporates a simple backward looking element (a stylized form of habit):

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

In Table 6 we report the results from estimating this variant of the model, presented as a grid search over the value of the adjustment parameter δ . The data seem to prefer a small but positive value of δ , which helps to account for the positive autocorrelation in the univariate process for consumption growth (see Table 1, line 2).

Our second extension is motivated by the discussion of labor hoarding and pro-cyclical

| δ | ρ | σ_u | σ_ϵ | σ_η | σ_ν | ML |
|----------|--------|------------|-------------------|---------------|--------------|--------|
| 0 | 0.8785 | 0.0068 | 0.0008 | 0.0063 | 0.0086 | 1073.3 |
| 0.1 | 0.8700 | 0.0071 | 0.0009 | 0.0066 | 0.0080 | 1075.9 |
| 0.2 | 0.8591 | 0.0075 | 0.0011 | 0.0070 | 0.0072 | 1074.8 |
| 0.3 | 0.8412 | 0.0082 | 0.0013 | 0.0075 | 0.0062 | 1068.8 |
| 0.4 | 0.7823 | 0.0092 | 0.0020 | 0.0081 | 0.0035 | 1057.0 |
| 0.5 | 0.6915 | 0.0107 | 0.0033 | 0.0089 | 0.0002 | 1044.4 |
| 0.6 | 0.7126 | 0.0130 | 0.0037 | 0.0110 | 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 | 0.6480 | 0.0567 | 0.0200 | 0.0456 | 0.0033 | 796.0 |

Table 6: Maximum Likelihood Estimation: Slow Consumption Adjustment

| α | ρ | σ_u | σ_ϵ | σ_η | σ_ν | ML |
|----------|--------|------------|-------------------|---------------|--------------|--------|
| 0 | 0.8910 | 0.0067 | 0.0007 | 0.0063 | 0.0089 | 1073.2 |
| 0.1 | 0.8989 | 0.0069 | 0.0007 | 0.0065 | 0.0067 | 1072.9 |
| 0.2 | 0.9110 | 0.0072 | 0.0006 | 0.0069 | 0.0052 | 1071.6 |
| 0.3 | 0.9249 | 0.0077 | 0.0006 | 0.0074 | 0.0039 | 1068.2 |
| 0.4 | 0.8948 | 0.0085 | 0.0009 | 0.0080 | 0.0 | 1064.6 |
| 0.6 | 0.9434 | 0.0114 | 0.0006 | 0.0111 | 0.0 | 1034.1 |
| 0.8 | 0.9645 | 0.0229 | 0.0008 | 0.0225 | 0.0 | 937.6 |
| 1 | 0.0070 | 0.0067 | 0.0067 | 0.0006 | 0.0857 | 391.1 |

Table 7: Maximum Likelihood Estimation: Labor hoarding

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 that 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.

To capture endogenous responses of productivity, we extend the model by assuming that the process a_t captures the exogenous component of productivity, while actual productivity, denoted by \tilde{a}_t , responds to increases in employment according to the relation:

$$\tilde{a}_t = a_t + \alpha(c_t - a_t).$$

Table 7 displays the Maximum Likelihood estimation for this case, as a grid over values for α . In this case, the model fits the data better with no endogenous productivity responses,

i.e., with $\alpha = 0$. However, the likelihood is relatively flat for low levels of α . Notice that, in that region, the model compensates for higher values of α by choosing lower estimates for σ_ν . To interpret this result, remember from Section 4.1 that higher values of σ_ν are associated to a higher coefficient of correlation between the innovations of consumption and productivity u_t^c and u_t^a . Allowing for endogenous productivity, gives us an alternative channel to explain this correlation. The results in Table 7 show that, having only data on consumption and observed productivity, it is hard to distinguish the role of these two channels.

6 Conclusions

On the methodological side, we have explored the problem of estimating models with news and noise—which we think provide an appealing description of the cycle. We have shown the limits of SVAR estimation, and shown how these models can be estimated with structural methods. This implies that to identify the role of news and noise in fluctuations one must rely more heavily on the model’s structure. In this paper, a central role for identification was played by the consumer’s Euler equation, that is, by the assumption that current movements in consumption are primarily driven by changes in the consumers’ expectations on the economy’s long run potential.

On the empirical side, the data appear quite consistent with a view of fluctuations where the pattern of technological change is smooth, subject to random shocks which only build up slowly, while most of the short-run action in consumption and output comes from noisy information on these long-run trends. Clearly, we need to extend the model in many dimensions before having confidence in these conclusions. In particular, adding investment seems an essential step in building models of the business cycle driven by anticipations.

Another natural extension is to add variables to the empirical exercise, to better capture consumers’ expectations about the future. For example, one could include financial market prices, following Beaudry and Portier (2006), or survey measures of consumer confidence, as Barsky and Sims (2008). However, the analysis in Section 3, where we allowed the econometrician to directly observe all the signals observed by the consumers, shows that adding variables is not sufficient, in general, to solve the identification problems of SVARs.

Finally, it is useful to notice that the applicability of SVAR methods depends crucially on the way in which one models the information structure. In models where the consumer exactly observes shocks which will affect productivity in the future, invertibility problems may be less damning (see our comments in Section 3.5 and the analysis in Sims (2009)). However, we think that, in many instances, signal extraction models provide a more realistic

and flexible description of the informational environment. When dealing with these models, the researcher can choose, depending on the question at hand, either to limit attention to the innovation representation of the consumers' forecasting problem or to take the structural approach developed here.

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

$$E \sum_{t=0}^{\infty} \beta^t U(C_t, N_t),$$

with

$$U(C_t, N_t) = \log C_t - \frac{1}{1+\zeta} N_t^{1+\zeta},$$

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

$$Y_{j,t} = A_t N_{j,t}. \tag{16}$$

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

$$Q_t B_{t+1} + \int_0^1 P_{j,t} C_{j,t} dj = B_t + W_t N_t + \int_0^1 \Pi_{j,t} dj, \tag{17}$$

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 (17) 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

$$i_t = i^* + \phi \pi_t, \tag{18}$$

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:

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

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

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

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

Appendix B. Consumers' Kalman filter

Define the matrices

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

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 $P \equiv \text{Var}_{t-1}[\xi_t]$. 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. Proof of Proposition 1

Let w_t be an identified shock, corresponding to a linear combination of current and past observables. Applying the law of iterated expectations we get

$$E \left[c_{t+k} | w_t, \mathcal{I}_{t-1}^e \right] = E \left[\lim_{j \rightarrow \infty} E \left[a_{t+k+j} | \mathcal{I}_{t+k} \right] | w_t, \mathcal{I}_{t-1}^e \right] = \lim_{j \rightarrow \infty} E \left[a_{t+j} | w_t, \mathcal{I}_{t-1}^e \right],$$

for all $k \geq 0$ and, similarly,

$$E [c_{t+k} | \mathcal{I}_{t-1}^e] = \lim_{j \rightarrow \infty} E [a_{t+j} | \mathcal{I}_{t-1}^e].$$

It follows that the response of consumption to w_t is constant and equal to

$$E [c_{t+k} | w_t, \mathcal{I}_{t-1}^e] - E [c_t | \mathcal{I}_{t-1}^e] = \lim_{j \rightarrow \infty} E [a_{t+j} | w_t, \mathcal{I}_{t-1}^e] - \lim_{j \rightarrow \infty} E [a_{t+j} | \mathcal{I}_{t-1}^e],$$

for all $k \geq 0$.

Appendix D. Econometrician's Kalman Filter

The econometrician's state vector is given by

$$\xi_t^E \equiv (x_t, x_{t-1}, z_t, x_{t|t}, x_{t-1|t}, z_{t|t})'.$$

Rewrite the dynamics of the vector of consumer expectations $(x_{t|t}, x_{t-1|t}, z_{t|t})$, from (9), as follows:

$$\begin{aligned} \begin{bmatrix} x_{t|t} \\ x_{t-1|t} \\ z_{t|t} \end{bmatrix} &= A \begin{bmatrix} x_{t-1|t-1} \\ x_{t-2|t-1} \\ z_{t-1|t-1} \end{bmatrix} + B \begin{bmatrix} 1 + \rho_x & -\rho_x & \rho_z \\ 1 + \rho_x & -\rho_x & 0 \end{bmatrix} \begin{bmatrix} x_{t-1} \\ x_{t-2} \\ z_{t-1} \end{bmatrix} + \\ &+ B \begin{bmatrix} 1 \\ 1 \end{bmatrix} \epsilon_t + B \begin{bmatrix} 1 \\ 0 \end{bmatrix} \eta_t + B \begin{bmatrix} 0 \\ 1 \end{bmatrix} \nu_t. \end{aligned}$$

Then the state ξ_t^E evolves according to:

$$\xi_t^E = Q \xi_{t-1}^E + R (\epsilon_t, \eta_t, \nu_t)'. \quad (19)$$

where the matrices Q and R are given by

$$Q = \begin{bmatrix} 1 + \rho_x & -\rho_x & 0 & & \\ & 1 & 0 & 0 & \mathbf{0} \\ & 0 & 0 & \rho_z & \\ B \begin{bmatrix} 1 + \rho_x & -\rho_x & \rho_z \\ 1 + \rho_x & -\rho_x & 0 \end{bmatrix} & & & & A \end{bmatrix},$$

References

- [1] Aguiar, M. and G. Gopinath, 2007, “Emerging Market Business Cycles: The Cycle is the Trend,” *Journal of Political Economy*, 115(1): 69-102.
- [2] Anderson, E. W., L. P. Hansen, E. R. McGrattan, and T. J. Sargent, 1996, “Mechanics of Forming and Estimating Dynamic Linear Economies,” in *Handbook of Computational Economics* ed. by H. M. Amman, D. A. Kendrick, and J. Rust, Vol. 1, Chap. 4, 171-252, Elsevier.
- [3] Barsky, R. and E. Sims, 2008, “Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence”, mimeo, Michigan.
- [4] Beaudry, P. and F. Portier, 2006, “Stock Prices, News, and Economic Fluctuations,” *American Economic Review*, 96(4): 1293-1307.
- [5] Blanchard, O. and D. Quah, 1989, “The Dynamic Effects of Aggregate Demand and Aggregate Supply Disturbances,” *American Economic Review*, 79(4): 654-673.
- [6] Blundell, R., Preston, I., 1998, “Consumption Inequality and Income Uncertainty,” *Quarterly Journal of Economics*, 113(2): 603-640.
- [7] Boz, E., C. Daude and C. B. Durdu, 2008, “Emerging Market Business Cycles Revisited: Learning About the Trend,” FRB International Finance Discussion Paper No. 927.
- [8] Christiano, L., C. Ilut, R. Motto and M. Rostagno, 2007, “Signals: Implications for Business Cycles and Monetary Policy,” mimeo, Northwestern University.
- [9] Fernández-Villaverde, J., J. F. Rubio, T. Sargent and M. Watson, 2007, “A, B, C, (and D)’s for Understanding VARs,” *American Economic Review*, 97(3), 1021-1026.
- [10] Galí, J., 1992, “How Well Does the IS-LM Model Fit Postwar US Data?”, *Quarterly Journal of Economics*, 107(2): 709-738.
- [11] Galí, J., 1999, “Technology, Employment and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?” *American Economic Review*, 89(1), 249-271.
- [12] Galí, J., 2008, *Monetary Policy, Inflation and the Business Cycle: An Introduction to the New Keynesian Framework*. Princeton University Press.

- [13] Jaimovich, N. and S. Rebelo, 2006, "Can News about the Future Drive the Business Cycle?" mimeo, Northwestern University.
- [14] Leeper, E. M., T. B. Walker and S. S. Yang, 2009, "Fiscal Foresight and Information Flows," NBER Working Paper 14630.
- [15] Lorenzoni, G., 2009, "A Theory of Demand Shocks," *American Economic Review*, forthcoming.
- [16] 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.
- [17] Quah, D., 1991, "The Relative Importance of Permanent and Transitory Components: Identification and Some Theoretical Bounds," NBER Technical Working Papers 106.
- [18] Rodriguez Mora, J., Schulstad, P., 2007, "The Effects of GNP announcements on Fluctuations of GNP Growth," *European Economic Review*, 51(8), 1922-1940.
- [19] Schmitt-Grohé, S. and M. Uribe, 2008, "What's News in Business Cycles," mimeo, Columbia University.
- [20] Shapiro, M.D. and M.W. Watson, 1988, "Sources of Business Cycle Fluctuations," *NBER Macroeconomics Annual*, MIT Press.
- [21] Sims, C. and T. Zha, 2006, "Does Monetary Policy Generate Recessions?" *Macroeconomic Dynamics*, 10(2), 231-272.
- [22] Sims, E., 2009, "Non-Invertibilities and Structural VARs," mimeo, Michigan.