Institute for Advanced Development Studies

09/2012

News, Noise, and Fluctuations: An Empirical Exploration

by:

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

Development Research Working Paper Series

09/2012

October 2011

The views expressed in the Development Research Working Paper Series are those of the authors and do not necessarily reflect those of the Institute for Advanced Development Studies. Copyrights belong to the authors. Papers may be downloaded for personal use only.

News, Noise, and Fluctuations: An Empirical Exploration

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

October 2012

Abstract

We explore empirically models of aggregate fluctuations with two basic ingredients: agents form anticipations about the future based on noisy sources of information and 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 agents' estimates of these fundamentals (noise). We use a simple forward-looking model of consumption to address some methodological issues: structural VARs cannot be used to identify news and noise shocks in the data, but identification is possible via a method of moments or maximum likelihood. Next, we use U.S. data to estimate both our simple model and a richer DSGE model with the same information structure. Our estimates suggest that noise shocks play an important role in short-run consumption fluctuations.

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

JEL Codes: E32, C32, D83

*IMF, MIT and NBER; EIEF; Northwestern University and NBER. We thank for useful comments Daron Acemoglu, Marios Angeletos, Ricardo Caballero, Larry Christiano, Martin Eichenbaum, Matteo Iacoviello, Cosmin Ilut, Lars Hansen, Marco Lippi, Anna Mikusheva, Roberto Rigobon, Paul Schrimpf, Chris Sims, Eric Sims, Raf Wouters, Mark Watson, Iván Werning and seminar participants at several institutions.

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.

This view 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.

In this paper, we examine whether this view is consistent with the data. We reach three main conclusions, the first two methodological, the third substantive:

Structural VARs typically cannot recover news and noise shocks. 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.

While structural estimation methods cannot recover the actual time series for news and noise shocks either, they can recover underlying structural parameters, and thus the relative role and dynamic effects of news and noise shocks.

Estimation of both a simple model, and then of a more elaborate DSGE model suggest that agents indeed solve such a signal extraction problem, and that noise shocks play an important role in determining short run dynamics.

Recent efforts to estimate business cycle models in which expectations about the future play an important role include Christiano, Ilut, Motto and Rostagno (2007) and Schmitt-Grohé and Uribe (2008). Those papers follow the approach of Jaimovich and Rebelo (2006) and model news as perfectly anticipated productivity changes that will occur at some future date. We share with those papers the emphasis on structural estimation. The main difference is the use of a signal extraction model for consumers' information and our focus on disentangling the role of news and noise.[1]

[1] Beaudry and Portier (2004) is an early example of a signal extraction model in the recent literature on

The paper is also related to recent papers that have pointed out that the way we model the agents' information structure may affect the applicability of structural VARs methods, e.g, Fernández-Villaverde, Rubio-Ramírez, Sargent and Watson (2007) and Leeper, Walker and Yang (2009).

The paper is organized as follows. Section 1 presents a simple analytical model around which the discussion is best organized. Section 2 discusses estimation. Section 3 presents the results of estimation of a larger DSGE model.

1 A simple model

We begin with 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.

Productivity is driven by two shocks: a permanent shock and a transitory shock.[2] Consumers do not observe the two shocks separately, but only the realized level of productivity. The permanent shock introduces uncertainty about the economy's long run fundamentals. The presence of the transitory shock implies that consumers cannot back out the permanent shock from productivity observations, thus creating a signal extraction problem.

Consumers have access to an additional source of information, as they observe a noisy signal of the permanent component of productivity. This adds a third source of fluctuations, a shock to the error term in the signal, which we call "noise shock."

Consumers solve their signal extraction problem, form expectations about future productivity, and choose spending based on these expectations. Because of nominal rigidities, spending determines output in the short run.

Now to the specific assumptions.

Productivity a_t (in logs) is the sum of two components, the permanent component x_t and the transitory component z_t,

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

business cycles driven by expectations about the future. Lippi and Neri (2007) have a signal extraction DSGE model and estimate it by maximum likelihood, but do not consider shocks to expectations about the future.

[2]Permanent shock is a slight (and common) misnomer, as it refers to a shock with permanent effects that build up gradually.

The permanent component follows the unit root process

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

The transitory component follows the stationary process

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

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.

For most of the paper, we assume that

$$\rho_x = \rho_z = \rho, \tag{4}$$

and that ρ and the variances σ_ϵ^2 and σ_η^2 satisfy the restriction

$$\rho \sigma_\epsilon^2 = (1 - \rho)^2 \sigma_\eta^2. \tag{5}$$

The motivation for these restrictions is that, together, they imply that the univariate process for a_t is a random walk, that is

$$E[a_{t+1} | a_t, a_{t-1}, ...] = a_t.$$

This random walk representation is analytically convenient and, as will be seen below, also broadly in line with actual productivity data.[3][4] To see why this property holds, note first that the implication is immediate when $\rho = \sigma_\eta = 0$. Consider next the case in which ρ is positive and both variances are positive. An agent who observes a productivity increase at time t, can attribute it to an ϵ shock and forecast future productivity growth, or to an η shock and forecast mean reversion. When (5) is satisfied, these two considerations exactly balance and expected future productivity is equal to current productivity.[5]

On top of observing the realized productivity level a_t each period, consumers receive

[3]See Table 1.

[4]A similar process (with full information) was recently used by Aguiar and Gopinath (2007) in an open economy calibration exercise. Boz, Daude and Durdu (2008) explore the role of different informational assumptions in that context.

[5]The proof is as follows. In general, (1)-(3) imply

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

a noisy signal about the permanent component x_t. The signal is given by

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

where ν_t is i.i.d. normal with variance σ_ν^2.

We assume that consumers set consumption (in logs) c_t equal to their long run productivity expectations

$$c_t = \lim_{j \to \infty} E_t[a_{t+j}], \tag{7}$$

where E_t is the expectation conditional on the consumers' information at date t, i.e., conditional on current and past values of a_t and s_t. We drastically simplify the determination of output by assuming that consumption is the only component of demand, and that output is fully determined by the demand side. Thus, output (in logs) is equal to consumption

$$y_t = c_t,$$

Finally, we assume a linear production function in labor, so that the labor input n_t adjusts to produce y_t, given the current productivity a_t, and

$$n_t = y_t - a_t.$$

In the Online Appendix 6.2 we show that this model is the limit case of a standard new Keynesian model with Calvo pricing and a simple inflation targeting rule, when the frequency of price adjustment goes to zero. A useful property of this simple model is that consumption, by construction, is a random walk:

$$c_t = E_t[c_{t+1}], \tag{8}$$

which simply follows from the law of iterated expectations.

---

and

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

If (5) holds, these yield $Var[\Delta a_t] > 0$ and $Cov[\Delta a_t, \Delta a_{t-j}] = 0$ for all $j > 0$. Quah (1990, 1991) offers general results on the decomposition of a univariate process in permanent and transitory components with orthogonal innovations.

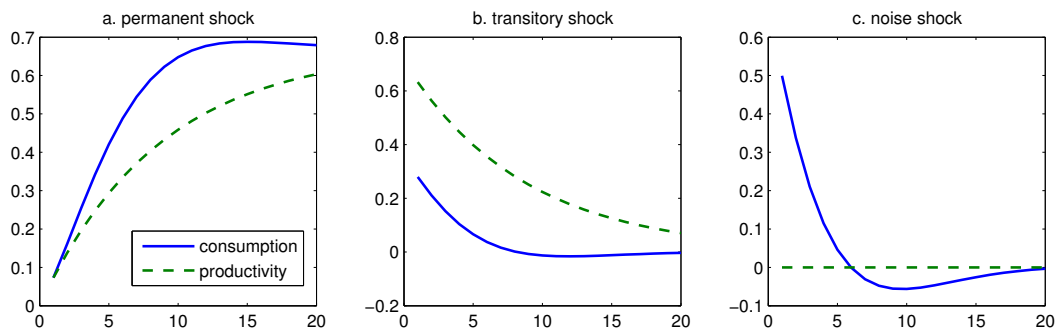

Figure 1: Impulse responses

1.1 Solving the model

The only endogenous variable in the model is c_t, and we now solve for it. Using (2) we can compute the expected value of cumulated productivity growth in the long run

$$\lim_{j\to\infty} E_t[x_{t+j} - x_t] = \frac{\rho}{1-\rho} E_t[x_t - x_{t-1}].$$

Since the transitory component disappears in the long run, we can replace a with x in (7) and, rearranging the equation above, get consumption:

$$c_t = \frac{1}{1-\rho}\left(E_t[x_t] - \rho E_t[x_{t-1}]\right). \tag{9}$$

To complete the solution of the model, one needs to solve the consumers' signal extraction problem to express the expectations of x_t and x_{t-1} in terms of current and lagged values of the shocks $(\epsilon_t, \eta_t, \nu_t)$. This is done using standard Kalman filtering. The resulting expressions, which are not particularly simple, are given in Appendix 5.1.

Figure 1 shows the responses of consumption and productivity to our three shocks. We use as parameters the estimated parameters from Section 2 below. The time unit is the quarter and the impulses are one standard deviation positive shocks. The persistence parameter is $\rho = 0.89$, implying slowly building permanent shocks and slowly decaying transitory shocks. The standard deviations of the two technology shocks are $\sigma_\epsilon = 0.07\%$ and $\sigma_\eta = 0.63\%$ and that of the noise shock is $\sigma_\nu = 0.89\%$, implying a fairly noisy signal.

In response to a permanent shock ϵ_t, productivity builds up slowly over time—the implication of a high ρ_x—and consumption also increases slowly. This reflects the fact that the volatilities of transitory and noise shocks are relatively large, so that it takes a while for consumers to recognize the permanent shock and adjust consumption. For

our parameter values, consumption—which depends on the expected value of long-run productivity—initially increases faster than productivity, generating a transitory increase in employment. A more volatile transitory shock or a less informative signal, can yield a slower consumption adjustment, generating an initial drop in employment.

In response to a transitory shock η_t, productivity initially increases, and then slowly declines over time. As agents put some weight on the productivity increase being due to a permanent shock, consumption initially increases. As agents learn that it was only a transitory shock, consumption returns back to normal. 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 noise shock ν_t, 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.

In the next section, we ask whether and how we can recover the responses in Figure 1 from the data.

2 Identification and estimation

We now turn to issues of identification and estimation.

First, we derive the reduced form VAR representation of the process for consumption and productivity and show that it is typically non-invertible. The result is more general than our model and implies that it is not possible to use simple semi-structural identification assumptions to estimate the economy's responses to shocks.

Second, we show how, in our simple model, identification of parameters can be achieved using three moments in the data. This exercise shows how information in the data can be used to shed light on the role of news and noise shocks.

Third, and in preparation for the estimation of the larger DSGE model in the next section, we show how the model can be estimated by maximum likelihood, thus using all the information in the data.

2.1 Reduced form VAR

Given our assumptions, the reduced form VAR representation for c_t and a_t takes the following simple form (this assumes that the econometrician does not observe s_t. We return to the issue below):

$$c_t = c_{t-1} + u_t^c, \tag{10}$$

$$a_t = \rho a_{t-1} + (1-\rho)\, c_{t-1} + u_t^a, \tag{11}$$

where u_t^c and u_t^a are innovations with respect to the econometrician information set.

Let us provide the steps behind (10) and (11). To derive the first it is sufficient to notice that consumption satisfies the random walk property (8), where the expectation is conditional on past values of a_t and c_t. To derive the second, use (1) and (3) to get

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

Next, use (8) and (9) and the law of iterated expectations to get

$$(1-\rho)c_{t-1} = (1-\rho)E_{t-1}[c_t] = E_{t-1}[x_t - \rho x_{t-1}].$$

Combining these two results yields

$$E_{t-1}[a_t - \rho a_{t-1} - (1-\rho)c_{t-1}] = 0,$$

which implies the representation (11).

The interesting feature of this representation is the presence of c_{t-1} in the productivity equation. Recall that the univariate representation of productivity is a random walk, by assumption. But when we move to a multivariate representation, past consumption helps to predict productivity. The reason is that consumption embeds the additional information on x_t that the consumers obtain from observing s_t.

2.2 Structural VAR

Suppose we run a reduced form VAR in (c_t, a_t) and obtain the reduced form innovations (u_t^c, u_t^a).

It is obvious that we cannot recover the original three shocks $(\epsilon_t, \eta_t, \nu_t)$ from two reduced form innovations. Only in two special cases can this be done:

The first is the case of perfect signal extraction, when $\sigma_\nu = 0$. In this case (10) and (11) 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. If we impose 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. So in this case a structural VAR approach works.

The second is the case of no signal extraction, when $\sigma_\nu \to \infty$. In this case, consumers only observe a_t and our random walk assumption implies that consumption and productivity are perfectly correlated with

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

where u_t denotes the common innovation in the two variables. In this case, it is not possible to recover ϵ_t and η_t from the single innovation u_t. But the decomposition between temporary and permanent shocks is now irrelevant, given that no information is available to separate them. We can then take the random walk representation of productivity as our primitive and interpret the productivity innovation as the single, permanent shock. In terms of this alternative representation, a structural VAR approach works.

Once we move away from these two cases however, and have a partially informative signal, the reduced form VAR representation is non-invertible and a structural VAR approach cannot be used.

This conclusion raises however two questions:

First, what if the econometrician also observes the signal, so that he can estimate a trivariate reduced form in (c_t, a_t, s_t), with residuals (u_t^c, u_t^a, u_t^s)? Even in this case, the answer remains the same. The reason is that agents' decisions are functions of their expectations, so even if the econometrician observes the three variables (c_t, a_t, s_t), the first variable is a function of the other two, which implies that there are only two independent innovations driving the system. It is still impossible to recover three orthogonal shocks from two innovations. Lemma 1 in the appendix, formulated in the context of a general signal extraction model, shows that singularity is endemic to this class of models.

Second, one might still hope that long-run identification restrictions can be used to separate the effect of the permanent shock ϵ_t from the *combined* effect of the other two shocks, η_t and ν_t. Unfortunately, this partial identification also fails, and in dramatic fashion. In our model, in which consumption follows a random walk, the following result holds: Consumption displays a flat impulse response to *any* identified shock, whether identified à la Blanchard and Quah (1989) or by any other short-term or long-term restriction. This result is proven formally as Proposition 1 in the Online Appendix (Section 6.1). The intuition is that since consumption is a random walk conditional on the consumers' information, an econometrician with access to the same information or less, cannot identify any shock with non-flat effects on consumption.

Figure 2 shows what happens through the results of a simulation. We generate data from our model—with the parameters used for Figure 1—and run a structural VAR with long-run restrictions à la Blanchard and Quah (1989) to identify a permanent and a temporary shock (BQ shocks from now on). The figure shows the estimated impulse responses to the two BQ shocks (dashed lines) and the impulse responses to the three original shocks in the model (solid lines). The two panels on the left focus on shocks with permanent effects. For both productivity and consumption, the BQ shock has larger effects on impact and less of a gradual build up in later periods, relative to the original shock ϵ.[6] This is especially pronounced for consumption, where the response to the BQ shock is virtually flat.[7] The two panels on the right show the responses to shocks with temporary effects.[8] The productivity response to the BQ transitory shock is the only one close to that of the original η shock. The identified response of consumption to the BQ transitory shock is zero, unlike those of η and ν.

2.3 Matching moments

While the lack of invertibility implies that structural estimation cannot recover the shocks themselves, the use of more model restrictions than in structural VARs allows us to recover the underlying parameters, and thus the dynamic effects of the shocks.

In our model three moments from the data are sufficient to identify all the parameters. First, ρ is identified using the reduced form relation (11). Given ρ, it is then easy to

[6] All the responses are to one standard deviation shocks.

[7] The theoretical result mentioned above—that the response of consumption to any identified shock is flat—holds only asymptotically, as the size of the sample goes to infinity. In short samples, the impulse responses are only approximately flat, as in Figure 2.

[8] The thick line corresponds to the η shock, the thin line to the ν shock, the dashed line to the BQ transitory shock.

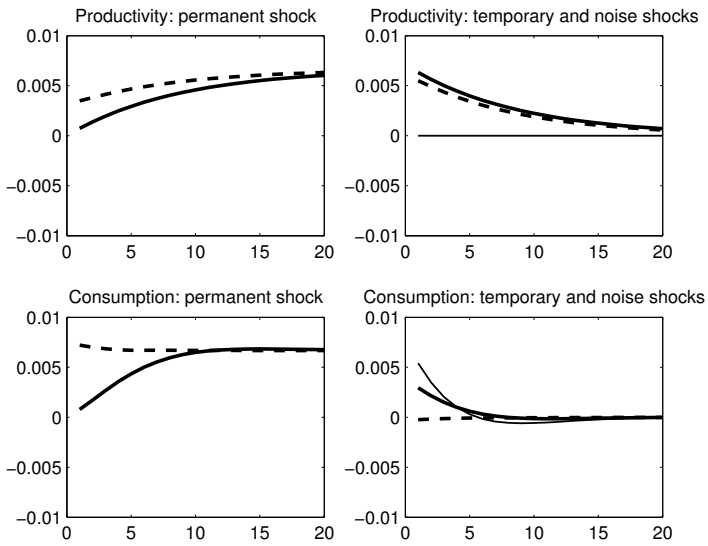

Figure 2: Model and SVAR-identified impulse responses

recover the variances σ_ϵ^2, σ_η^2 and σ_ν^2. This identification exploits the model's assumptions on consumers' forward looking behavior and rational expectations.[9]

We now go through each step of our moment-based estimation using U.S. quarterly data. This will also allow us to show that our reduced form benchmark model (10)-(11) fits the time series facts for productivity and consumption fairly well.

We measure productivity a_t as the logarithm of the ratio of GDP to employment and consumption c_t as the logarithm of the ratio of NIPA consumption to population. Our sample is 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 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 and remove it from the consumption series.[10]

Some basic features of the time series for productivity and consumption are presented

[9]The use of the permanent income logic together with rational expectations to identify temporary and permanent shocks connects our approach to a large body of work on household income dynamics, e.g., Blundell and Preston (1998).

[10]In the context of our approach, where we are trying to isolate potentially low frequency movements in productivity, this is an imperfect solution. But, given our purposes, it seems a reasonable first pass assumption.

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 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.

The first step of our identification uses the reduced form equation (11) to recover ρ. Writing (11) as a cointegrating regression, we have

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

which can be estimated by OLS. Our estimate is reported on Line 3 of Table 1. 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.

The model provides alternative ways of estimating ρ, exploiting the correlation of productivity growth and consumption at different horizons. Namely, the model implies

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

for any $j \geq 0$, where $u_t^{a,j}$ is a disturbance uncorrelated to the econometrician's information at date t.[11] Lines 5 to 7 explore this implication.[12] The results are roughly consistent with the model predictions, and all point to relatively high values for ρ (reported in the last column). The idea that the forecasting power of the consumption-to-productivity ratio tells us something about consumers' information about future productivity is closely related to a similar observation made by Cochrane (1994) in terms of the consumption-to-output ratio. Indeed this observation was the motivating reason for Cochrane's (1994) early suggestion to introduce news shocks in business cycle models.

The second step of our identification is to estimate σ_ϵ and σ_η. For this, we exploit

[11] This equation is obtained by induction. It holds for $j = 0$ from (11). If it holds for j, then $E_t\left[a_{t+j}\right] = \left(1-\rho^j\right) c_t + \rho^j a_t$. 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 (10) and (11), the third from rearranging.

[12] The standard errors are corrected for the presence of autocorrelation due to overlapping intervals using the Newey-West estimator.

	Dependent variable:	$\Delta a(-1)$	$\Delta c(-1)$	$(c-a)(-1)$	implied ρ
1	Δa	-0.06 (0.09)			
2	Δc		0.24 (0.08)		
3	Δa			0.05 (0.03)	0.95
4	Δa	-0.21 (0.10)	0.32 (0.12)	0.03 (0.02)	
5	$\Delta(8)a$			0.03 (0.15)	0.99
6	$\Delta(20)a$			0.31 (0.30)	0.98
7	$\Delta(40)a$			0.98 (0.43)	0.91

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 Newey-West window with 10 lags.

our univariate random walk assumption for a_t, that is, condition (5), which implies the following relations between the two variances and the variance of Δa_t:

$$\begin{aligned} \sigma_\epsilon^2 &= Var[\Delta a_t]/(1-\rho)^2 \\ \sigma_\eta^2 &= Var[\Delta a_t]/\rho \end{aligned}$$

Given a sample standard deviation of Δa_t equal to 0.67% and given our estimate for ρ, we get estimates $\sigma_\epsilon = 0.03\%$ and $\sigma_\eta = 0.65\%$. 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 third and last step is to recover the variance of the noise shock σ_ν. For this, we match the coefficient of correlation between the residual of regression (12) and consumption growth Δc. Notice that if the signal is perfectly informative ($\sigma_\nu = 0$) this correlation is positive but smaller than 1, while if the signal is completely uninformative ($\sigma_\nu \to \infty$) this correlation is 1.[13] Moreover, numerical results show that, given all other parameters, the coefficient of correlation is an increasing function of σ_ν. To get some intuition for this relation, notice that consumption is driven by the expected permanent component of productivity, while productivity itself is also driven by the temporary component. When the signal is more precise consumers can better separate the two components and so the innovations in consumption and productivity are less correlated. In the data, the coeffi-

[13]These bounds can be derived from the analysis in Section 2.1. 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.

cient of correlation is equal to 0.52, which is an intermediate value between the case of a perfectly informative signal (correlation 0.05) and the case of an uninformative signal (correlation 1). Therefore, the data point to the presence of a significant signal extraction problem, with an estimated standard deviation of the noise shock equal to $\sigma_\nu = 2.1\%$.

2.4 Maximum likelihood

We now turn to estimation by maximum likelihood. Conditional on the model being correctly specified, a maximum likelihood approach dominates the moment matching approach of the last section, as it fully incorporates all the restrictions implied by the model. For example, the maximum likelihood approach fully exploits the correlation between u_t^c and u_t^a implied by the model.

A maximum likelihood approach has the advantage that it can easily be extended to richer models, like the DSGE model of Section 3. In Appendix 5.1, we show how to compute the likelihood function for a general representative-agent model with signal extraction. The main idea is first to solve the consumer's Kalman filter to obtain the dynamics of consumer's expectations, as discussed in Section 1.1, and next to build the econometrician's Kalman filter, including in the list of unobservable state variables the consumer's expectations. This way of computing the likelihood function can also be used to apply Bayesian methods, as we shall do in Section 3.

Table 2 shows the results of estimation of the benchmark model presented as a grid over values of ρ from 0 to 0.99.[14] 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 $\rho = 0.89$, yielding the parameters reported on line 6.

Recall that the maximum likelihood approach uses all the implicit restrictions imposed by the model. This explains the difference between the estimates obtained with the estimated obtained by ML and those obtained by moment matching in Section 2.3. 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 closer to those in Section 2.3 and the likelihood gain from line 8 to line 6 is not very 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.

[14]For maximum likelihood estimation we used Dynare. 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).

	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

Note that the random walk assumption for productivity is not necessary for identification of the model's parameters. In particular, we can relax assumptions (4)-(5), allowing for different coefficients ρ_x and ρ_z in equations (2) and (3) and estimating independently σ_η and σ_ϵ. 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

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 were already given in Figure 1 of Section 1.1: 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.

Table 4 presents the implications of the estimated parameters for variance decomposition, showing the contribution of the three shocks to forecast error variance at different horizons. Noise shocks are the major source of short run volatility here, accounting for more than 70% of consumption volatility at a 1-quarter horizon and more than 50% at

Quarter	Perm. tech.	Trans. Tech.	Noise
1	0.016	0.235	0.749
4	0.269	0.198	0.533
8	0.683	0.087	0.229
12	0.832	0.046	0.122

Table 4: Variance decomposition

a one year horizon, 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. Clearly, the variance-decomposition implications are very sharp here because the baseline model only allows for three shocks. In the next section, we will see that noise shocks remain an important source of short-run consumption volatility in richer specifications that allow for more observables and more shocks.

At this stage, it is useful to compare the exercise here to traditional SVAR exercises, such as Shapiro and Watson (1989) and Gali (1992), that also use a small number of shocks and follow Blanchard and Quah (1988) to identify supply shocks. In those papers, transitory demand shocks typically explain a smaller fraction of aggregate volatility than our noise shock, and the permanent technology shock plays a bigger role. The analysis in Section 2.2 helps to explain the difference with our results, by showing that, asymptotically, a SVAR is biased towards assigning 100% of consumption volatility to the permanent shock.

2.5 Recovering states and shocks

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 shocks can be recovered from structural estimation?

Here the idea is to exploit the fact that the econometrician has access to the whole sample. Looking at what happens to productivity after the fact, we may be able to get a better sense of what the states and shocks were. In other words, using data from times 1 to T, we can form our best estimates of states and shocks at any time $t \leq T$. This is precisely the job of the Kalman smoother.

The top panel of Figure 3 plots estimates for the permanent component of productivity x_t obtained from our benchmark model. The solid line correspond to $x_{t|T}$, the econometrician's smoothed estimate of x_t. The dashed line is $x_{(t|t)|T}$, the econometrician's smoothed

estimate of the consumers' real time estimate of the same variable.

Looking 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 produce short-run volatility.

To gauge the short-run effects of expectational errors, 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 3, we plot the smoothed series for the consumers' real time expectations regarding long-run productivity, $x_{(t+\infty|t)|T} = (x_{(t|t)|T} - \rho x_{(t-1|t)|T})/(1-\rho)$, and compare it to the same expression computed using $x_{t|T}$ and $x_{t-1|T}$. The model generates large short-run consumption volatility out of temporary changes in consumers' expectations. 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. When we repeated the exercise using the full DSGE model of the next section (which allows for monetary policy shocks) this effect goes away.

The Kalman smoother also tell us what is the root-mean-square error (RMSE) of the estimates of x_t made both by the econometrician and by the consumer. It turns out that in steady state these two estimates coincide and the RMSE is 0.44% for estimates using data up to date t. If we can use all possible future data the RMSE halves, to 0.28%, but remains positive.[15] Appendix 6.4 contains more details.

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, the Kalman smoother gives estimates of ϵ_t, η_t, and ν_t using the entire time series available. Figure 4 plots these estimates for our benchmark model.

Notice the apparent high degree of autocorrelation of the estimated permanent shocks in the top panel of Figure 4. 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

[15] That is, this is the RMSE of the estimate of x_t based on data up to time $T > t$ when we let $T \to \infty$.

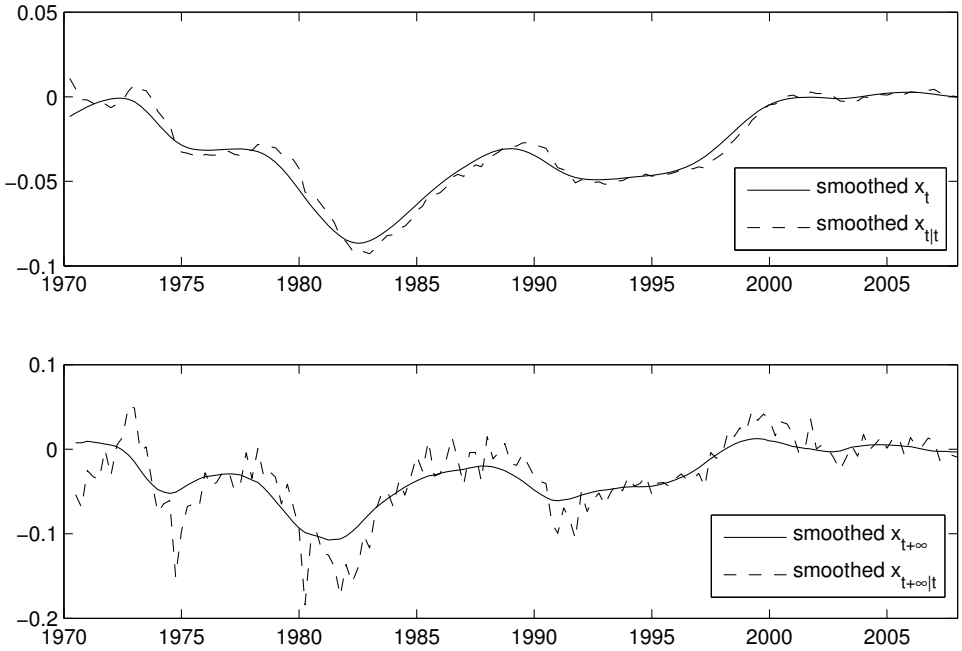

Figure 3: Smoothed estimates of the permanent component of productivity, of long-run productivity, and of consumers' real time expectations
Top panel: smoothed estimates of x_t (solid line) and $x_{t|t}$ (dashed line)
Bottom panel: smoothed estimates of $(x_t - \rho x_{t-1})/(1-\rho)$ (solid line) and $(x_{t|t} - \rho x_{t-1|t})/(1-\rho)$ (dashed line)

an observed permanent change in productivity.[16] 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.

3 A DSGE exercise

In this section, we start from the same productivity process and information structure of Section 1, but embed them in a small scale DSGE model. The model includes investment and capital accumulation, an explicit treatment of nominal rigidities and a monetary policy rule à la Taylor. The model also allows for variable capacity utilization and includes

[16] In Appendix 6.4 we show that the RMSE for the ϵ shock is very high, about 94% of the prior standard deviation σ_ϵ.

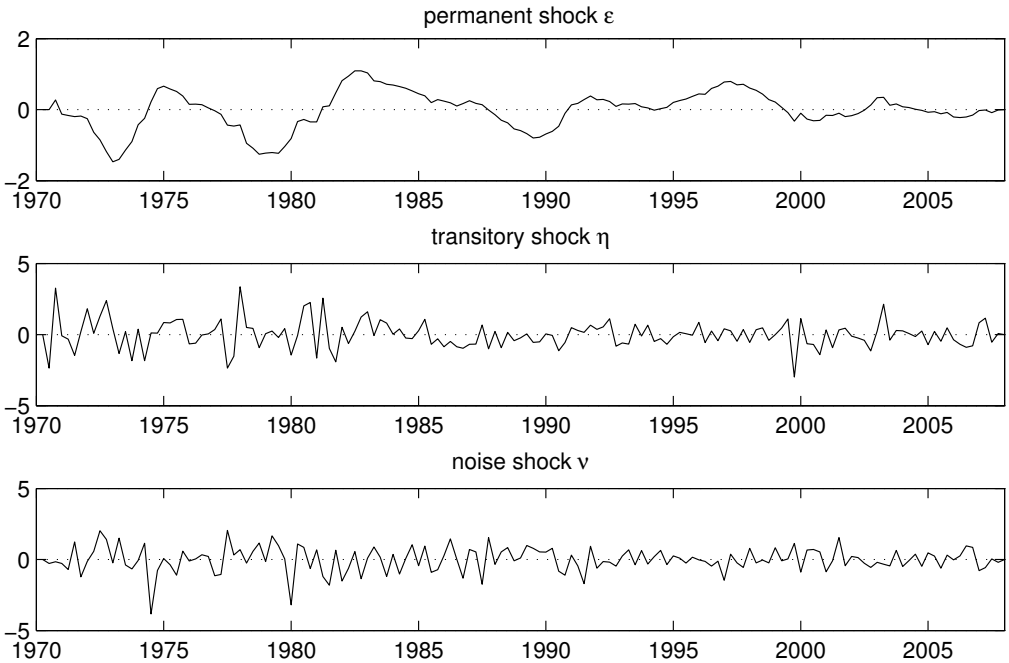

Figure 4: Smoothed estimates of the shocks

adjustment costs in consumption (habit) and investment, all elements that have been proposed in the literature to better capture the observed dynamics of aggregate quantities.

We have three objectives. First, we want to explore the robustness of our findings to a richer model, with a larger number of shocks. Second, we want to look at the response of investment to noise shocks. And finally we want to show that it is easy to estimate a DSGE model with a signal extraction information structure.

The model is estimated using Bayesian methods, as is now common for DSGE models with a relatively large number of parameters. The approach to compute the likelihood function is the one outlined in Section 2.4. However, a useful result makes the estimation easier. Namely, Lemma 2 in the Appendix shows that the model's information structure is observationally equivalent to the information structure of a model with full information and *correlated* shocks.[17] One can then estimate the full-information model subject to a restriction on the shocks' correlation matrix and, at the end, recover the parameters of the original signal extraction model.

[17]Therefore, a signal extraction model can be seen as a way of imposing restrictions on a class of models with correlated shocks.

3.1 Model

Since the model is standard, we describe here its main ingredients and leave the details and the log-linearization to the Online Appendix (Section 6.2). The model is similar to those in Smets and Wouters (2003) and Christiano, Eichenbaum, and Evans (2005). The preferences of the representative household are given by the utility function

$$E\left[\sum_{t=0}^{\infty}\beta^{t}\left(\log\left(C_{t}-hC_{t-1}\right)-\frac{1}{1+\zeta}\int_{0}^{1}N_{jt}^{1+\zeta}dj\right)\right], \tag{13}$$

where C_t is consumption, the term hC_{t-1} captures internal habit formation, and N_{jt} is the supply of specialized labor of type j. The presence of differentiated labor introduces monopolistic competition in wage setting as in Erceg, Henderson and Levin (2000). The capital stock $\bar{K}_t$ is owned and rented by the representative household and the capital accumulation equation is

$$\bar{K}_{t}=(1-\delta)\,\bar{K}_{t-1}+D_{t}\left[1-\mathcal{G}\left(I_{t}/I_{t-1}\right)\right]I_{t}, \tag{14}$$

where δ is the depreciation rate, D_t is a stochastic investment-specific technology parameter, and $\mathcal{G}$ is a quadratic adjustment cost in investment

$$\mathcal{G}(I_t/I_{t-1})=\chi(I_t/I_{t-1}-\Gamma)^2/2,$$

where Γ is the long-run gross growth rate of TFP. The model features variable capacity utilization: the capital services supplied by the capital stock $\bar{K}_{t-1}$ are

$$K_{t}=U_{t}\bar{K}_{t-1}, \tag{15}$$

where U_t is the degree of capital utilization and the cost of capacity utilization, in terms of current production, is $\mathcal{C}(U_t)\bar{K}_{t-1}$, where $\mathcal{C}\left(U_t\right)=U_t^{1+\xi}/\left(1+\xi\right)$.

The final good is a Dixit-Stiglitz aggregate of a continuum of intermediate goods, produced by monopolistic competitive firms, with staggered price setting à la Calvo (1983). Similarly, specialized labor services are supplied under monopolistic competition, with staggered nominal wages. The monetary authority sets the nominal interest rate following a standard inertial Taylor rule.

The model is estimated on U.S. time series for GDP, consumption, investment, employment, the federal funds rate, inflation, and wages, for the period 1954:3-2011:1. More

Table 5: Full DSGE: estimated parameters

	Parameter	Prior	Posterior	Conf. bands		Distribution	Prior st. dev.
h	Habit	0.5	0.5262	0.4894	0.5787	Beta	0.1
α	Production function	0.3	0.1859	0.1748	0.1933	Normal	0.05
ζ	Inv. Frisch elasticity	2	2.0871	1.0571	3.3012	Gamma	0.75
ξ	Capacity cost	5	3.4919	2.8912	4.3021	Normal	1
χ	Adjustment cost	4	4.3311	3.6751	5.5079	Gamma	1
θ	Calvo prices	0.66	0.8770	0.8545	0.8998	Beta	0.1
θ_w	Calvo wages	0.66	0.8690	0.8227	0.9183	Beta	0.1
γ_π	Taylor rule inflation	1.5	1.0137	1.0102	1.0568	Normal	0.3
γ_y	Taylor rule output	0.02	0.02	0.015	0.024	Normal	0.02
	Shock processes						
	Neutral technology and noise						
ρ		0.6	0.9426	0.9230	0.9618	Beta	0.2
σ_u		0.5	1.1977	1.0960	1.2975	Inv. Gamma	1
σ_ν		1	1.4738	0.7908	2.3176	Inv. Gamma	1
	Investment-specific						
ρ_d		0.6	0.4641	0.3263	0.5743	Beta	0.2
σ_d		0.15	11.098	8.4323	14.910	Inv. Gamma	1.5
	Mark ups						
ρ_p		0.6	0.7722	0.6991	0.8461	Beta	0.2
ϕ_p		0.5	0.4953	0.3749	0.6557	Beta	0.2
σ_p		0.15	0.1778	0.1508	0.2027	Inv. Gamma	1
ρ_w		0.6	0.9530	0.9534	0.9650	Beta	0.2
ϕ_w		0.5	0.5583	0.5125	0.6224	Beta	0.2
σ_w		0.15	0.3057	0.2847	0.3264	Inv. Gamma	1
	Policy						
ρ_q		0.4	0.0413	0.0024	0.0807	Beta	0.2
σ_q		0.15	0.3500	0.3148	0.3782	Inv. Gamma	1
ρ_g		0.6	0.9972	0.9938	0.9998	Beta	0.2
σ_g		0.5	0.2877	0.2680	0.3078	Inv. Gamma	1

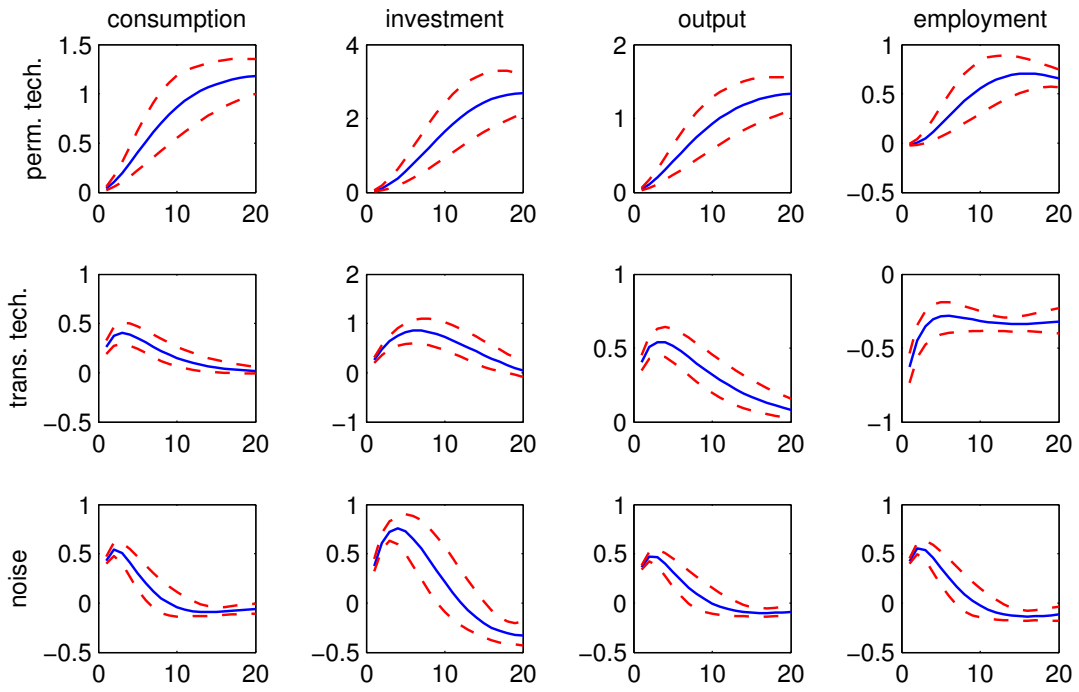

Figure 5: Impulse Responses, Bayesian DSGE, quantities

details on the data are in the Online Appendix (Section 6.3).

3.2 Results

The parameter estimates are reported in Table 5. Figures 5 and 6 show the impulse responses for our seven observed variables following the three shocks that are the focus of this paper: the permanent and transitory technology shocks, and the noise shock. Table 3.2 shows variance decomposition results for consumption, investment, and output, showing the contribution of the eight shocks in the model at different horizons.

First, let us look at the results for consumption.

Looking at the impulse responses in Figure 5, the responses of consumption to our three shocks are qualitatively similar to those shown in Figure 1 for the simple model of Section 1: in the short run consumption responds mostly to the noise and transitory technology shocks; the response to the noise shock dies down faster; the response to the permanent shock is small in the short run and builds up gradually. The main qualitative difference is a slightly hump-shaped response to the transitory and noise shocks, due to

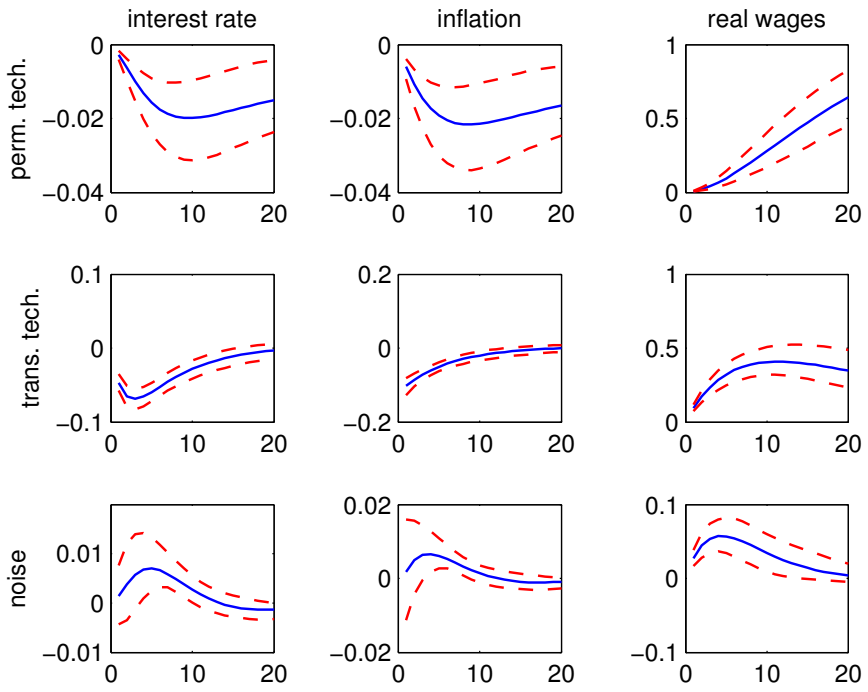

Figure 6: Impulse Responses, Bayesian DSGE, prices

the habit in preferences. The main quantitative differences are that the DSGE favors a larger coefficient of autocorrelation for growth shocks—$\rho = 0.94$ vs $\rho = 0.89$ in Section 2.4—and that it attributes larger volatility to both fundamental and noise shocks.

Looking at the variance decomposition in Table 3.2, we find that the noise shock is the main short-run driver of consumption, accounting for more than half of consumption volatility in the very short term and about 1/4th of it at a two year horizon. Relative to the simple model of Section 1, a sizeable fraction of short-run consumption volatility is now explained by the price markup shock.

The main novelty of the DSGE model, relative to our simple model earlier, is the presence of investment. The second column of Figure 5 shows that investment increases gradually and permanently after a permanent shock and has a hump-shaped response to a transitory shock. Following a noise shock the investment response is first positive and hump-shaped and later turns negative. What is happening is that at some point agents realize that the shock was just noise and the economy reverts to its original capital stock.

Two mechanisms drive up investment in the short run, following a noise shock. First, a higher expected marginal product of capital in the future leads to expected high future

Table 6: Variance decomposition

Consumption								
Quarter	Perm. tech.	Trans. Tech.	Noise	Inv. specific	Price markup	Wage markup	Monetary	Fiscal
1	0.004	0.186	0.512	0.001	0.205	0.037	0.001	0.055
4	0.064	0.246	0.430	0.002	0.117	0.039	0.006	0.095
8	0.331	0.198	0.245	0.003	0.063	0.024	0.015	0.121
12	0.577	0.117	0.134	0.003	0.034	0.013	0.017	0.106
Investment								
Quarter	Perm. tech.	Trans. Tech.	Noise	Inv. specific	Price markup	Wage markup	Monetary	Fiscal
1	0.000	0.005	0.011	0.971	0.006	0.006	0.000	0.000
4	0.003	0.017	0.021	0.936	0.008	0.016	0.000	0.000
8	0.031	0.036	0.027	0.869	0.009	0.027	0.000	0.001
12	0.120	0.046	0.025	0.769	0.009	0.029	0.000	0.003
Output								
Quarter	Perm. tech.	Trans. Tech.	Noise	Inv. specific	Price markup	Wage markup	Monetary	Fiscal
1	0.003	0.249	0.200	0.372	0.083	0.026	0.001	0.066
4	0.040	0.272	0.198	0.363	0.057	0.039	0.003	0.028
8	0.228	0.270	0.134	0.267	0.036	0.035	0.006	0.024
12	0.477	0.200	0.083	0.167	0.023	0.023	0.008	0.020

investment. This, combined with adjustment costs, leads to an increase in investment today.[18] Second, the short-run increase in consumption increases the expected marginal profitability of capital in the near term, with a direct effect on investment today.

In terms of variance decomposition, the noise shock only accounts for a small fraction of investment volatility, as virtually all short-run investment volatility is due to the investment-specific shock. To understand the relation between this result and the impulse responses in Figure 5, notice that a positive noise shock produces an increase in consumption and investment of similar magnitudes. However, unconditional investment volatility is larger than consumption volatility, so the investment-specific shock is needed to account for this extra volatility.

The responses of aggregate output follow from those of consumption and investment. In particular, in terms of variance decomposition the three most important drivers of output are the investment-specific shock, the transitory technology shock, and the noise shock, with the latter explaining about 20% of volatility at a 1-year horizon.

[18]The higher investment in the future is not eventually realized, as agents later learn that the shock was noise.

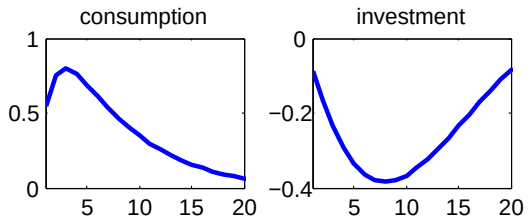

Figure 7: Impulse Responses to an Intertemporal Preference Shock

The DSGE model exploits the rich shock structure available and uses different shocks to explain separately the dynamics of consumption and investment. This is a common feature in estimated DSGE exercises. For example, in Justiniano, Primiceri and Tambalotti (2010) an investment-specific technology shock explains the largest fraction of investment volatility, while the largest fraction of consumption volatility is explained by a shock to intertemporal preferences. Our noise shock plays a role similar to an intertemporal-preference shock, as both appear as error terms on the right-hand side of the Euler equation.

We find expectation-based shocks like our noise shock more appealing than preference shocks, both on a priori grounds and because they impose more testable restrictions on consumption volatility. Moreover, when the model is estimated, noise shocks have an additional advantage over intertemporal preference shocks: they generate comovement of investment, consumption, and hours. That is, they produce aggregate responses in line with a standard definition of a business cycle.[19] To illustrate this difference, we have estimated our model removing noise—that is, under perfect information—and allowing for a standard autocorrelated shock to the consumers' discount factor. In Figure 7, we plot the responses of consumption and investment to a preference shock in this alternative estimation. Figure 7 shows that this preference shock produces negative comovement of consumption and investment.[20]

We conclude this section by observing that nominal rigidities and the monetary policy rule play an important role in producing substantial consumption volatility from noise shocks. Since actual productivity is unaffected when the noise shock hits, output increases above its natural level, generating inflation (bottom middle panel of Figure 6) and

[19]There is now a growing literature on the ability of various types of shocks to generate comovement. See Lorenzoni (2011) for a review.

[20]We do the estimations separately instead of combining noise and preference shocks in the same model, because having both shocks poses serious identification problems, making the maximization step unstable and the estimates very sensitive to the prior.

the central bank responds by raising interest rates (bottom left panel of Figure 6). The associated increase in real interest rates tends to dampen the consumption response. Sufficiently rigid prices and a sufficiently unresponsive Taylor rule imply that this dampening effect is not too strong. From experimenting with various combinations of parameters, we have reached the following conclusions. First, absent nominal rigidities the effects of a noise shock on consumption are muted and the effects on investment are reversed. Second, re-estimating the model imposing smaller values of the Calvo parameters (e.g., $\theta = \theta_w = 0.75$) or imposing a more responsive monetary policy rule (e.g., $\gamma_\pi = 1.5$), the simulated responses are close to the baseline. Third, if we impose at the same time smaller Calvo parameters and a more responsive policy rule, the effect of noise shocks on consumption can be considerably dampened.

Notice that our model uses log utility in consumption, as is common in the DSGE literature. This sets the intertemporal elasticity of substitution to one, making consumption highly sensitive to the real interest rate. It is quite possible that, with a lower elasticity of substitution, a more active monetary policy rule with more flexible prices would be less of a dampener of noise-driven consumption movements. We leave to future research the estimation of models with intertemporal elasticity different from one.[21]

Finally, notice that the recovery of states and shocks can be done in the DSGE model exactly as we did in Section 2.5 for the simple model, but, for reasons of space, the results are omitted.

4 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 business cycles. 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. Our simple model shows that a central role for identification is played by the consumer's Euler equation, which embeds the idea that consumption can be driven by changes in the consumers' long-run expectations. Our likelihood-based estimation exercises in Sections 2.4 and 3 show that signal extraction models can be easily estimated adapting common structural methods.

[21] We need preferences consistent with balanced growth, since we have a non-stationary technology process. Therefore, moving away from log utility will require to introduce non-separable preferences à la King, Plosser, and Rebelo (1988).

On the empirical side, the data appear consistent with a view of fluctuations where the pattern of technological change is smooth, subject to random shocks which only build up slowly, while a sizable fraction of short-run volatility in consumption and output comes from noisy information on these long-run trends.

A useful extension for future work is to add to the empirical exercise variables that capture directly information on consumers' expectations. For example, one could include financial market prices, following Beaudry and Portier (2006), or survey measures of consumer confidence, as in Barsky and Sims (2008). The analysis in Section 2.2, where we allow the econometrician to directly observe all the signals observed by the consumers, shows that adding these variables will not solve the identification problems of SVARs. But these variables can feed additional information into structural exercises like those of Sections 2.4 and 3, offering better ways of separating expectational shocks from other type of disturbances.

5 Appendix

In this appendix, we formulate a general representative agent dynamic linear model with signal extraction. Both the simple model of Section 1 and the full DSGE model of Section 3 are special cases of this formulation. We use this model for various purposes: (i) to set up the agents' Kalman filter used in the model solution (Section 1.1) and the econometrician's Kalman filter used to construct the likelihood function in Section 2.4; (ii) to derive the general singularity result for signal extraction models discussed in Section 2.2; (iii) to derive the equivalent full information model which simplifies estimation in Section 3.

Uncertainty is captured by the exogenous state vector X_t that follows the process

$$X_t = AX_{t-1} + BV_t, \tag{16}$$

where V_t is an n-dimensional vector of mutually independent i.i.d. shocks, with positive variance. The representative agent observes the m-dimensional vector

$$S_t = CX_t + DV_t. \tag{17}$$

In Sections 1 and 3, the state vector is $X_t = (x_t, x_{t-1}, z_t)'$, the shock vector is $V_t = (\epsilon_t, \eta_t, \nu_t)'$ and the vector of consumer observations is $S_t = (a_t, s_t)$. So the matrices A, B, C, D are

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

Let Y_t denote a vector of endogenous state variables controlled by the agent. Suppose the economic model can be described in terms of the stochastic difference equation

$$FE_t\left[Y_{t+1}\right] + GY_t + HY_{t-1} + MS_t + NE_t\left[S_{t+1}\right] = 0, \tag{18}$$

where F, G, H, M, N are matrices of parameters. Notice that the unobservable exogenous state X_t only enters the equilibrium through the observable vector S_t, reflecting the assumption that the information set of the representative agent is given only by past and current values of S_t and of the endogenous state Y_t. Suppose there is a unique stable solution of the model:

$$Y_t = PY_{t-1} + QS_t + RX_{t|t}, \tag{19}$$

where we use the notation $X_{t|t}$ for the agents' expectation $E\left[X_t | S_t, S_{t-1}, ...\right]$. The matrices

P, Q, R can be found solving the three matrix equations

$$FP^2 + GP + H = 0, \qquad (FP + G)\, Q + M = 0,$$
$$(FP + G)\, R + [F\,(QC + R) + NC]\, A = 0.$$

See Uhlig (1995) for techniques to solve the first equation in P. The solution of the other two is straightforward as they are linear in Q and R.

The economic model of Section 1 is given directly in the form (19), by equation (9). The economic model of Section 3 is presented in the Online Appendix.

5.1 Kalman filters

We can use the Kalman filter to express the agents' expectations $X_{t|t}$ in recursive form as

$$\begin{aligned} X_{t|t} &= AX_{t-1|t-1} + K(S_t - S_{t|t-1}) \\ &= (I - KC)AX_{t-1|t-1} + KS_t, \end{aligned} \tag{20}$$

where the matrix of Kalman gains K depends on the parameters of the productivity process. We assume that (20) is stable, i.e., all eigenvalues of $(I - KC)A$ are smaller than one in absolute value. Notice that stability of the filter does not require X_t to be stationary, e.g., the model used in Sections 1 and 3 is non-stationary and yet the filter is stable.

The vector of states for the econometrician is given by $(X_t, X_{t|t}, Y_t)$. The dynamics of X_t are given by (16). The dynamics of $X_{t|t}$ are given by

$$X_{t|t} = (I - KC)AX_{t-1|t-1} + KCAX_{t-1} + (KCB + KD)V_t,$$

which follows from (16) and (20). The dynamics of Y_t are given by (19). To set up the econometrician's Kalman filter we use the dynamic system just described for $\left(X_t, X_{t|t}, Y_t\right)$ and the observation equation

$$S_t^E = T \begin{bmatrix} Y_t & S_t \end{bmatrix}. \tag{21}$$

5.2 Singularity

Solving (19) backward and substituting in (21), we can express the econometrician's observables S_t^E in terms of distributed lags of the agents' observables S_t and of the agents' expectations $X_{t|t}$:

$$S_t^E = \Xi(L)(\ S_t \quad X_{t|t}\ )', \tag{22}$$

Define the vector of innovations for the econometrician as

$$U_t = S_t^E - E[S_t^E | S_{t-1}^E, S_{t-2}^E, ...]. \tag{23}$$

We say that the VAR in Y_t is invertible if V_t can be expressed as a linear combination of current and past values of U_t.

Lemma 1 *If the dimension of the agent's observation vector is smaller than the dimension of the shock vector, $m < n$, then the VAR in S_t^E is not invertible.*

Proof. The agents' Kalman filter can be solved backward to express $X_{t|t}$ as a function of current and past values of S_t. This, combined with (22) and (23), implies that S_t^E and thus U_t can be expressed as a function of current and past values of S_t. This implies that $Var[V_t|U_t, U_{t-1}, ...] \geq Var[V_t|S_t, S_{t-1}, ...]$. Standard derivations allow us to express the innovations in S_t, as $S_t - E[S_t|S_{t-1}, S_{t-2}, ...] = \Psi(L) V_t$. Since Ψ is $m \times n$, V_t cannot be expressed in terms of the agent's innovations, so $Var[V_t|S_t, S_{t-1}, ...] > 0$. Combining this with the inequality above yields $Var[V_t|U_t, U_{t-1}, ...] > 0$. ■

The main point of the lemma is that what matters is not the number of variables observed by the econometrician, but the number of variables observed by the agent. If the agent has not enough information to back up the shocks V_t, an econometrician cannot generate additional information on these shocks by observing the agent's behavior.

5.3 Equivalent full information model

Write the joint dynamics of $X_{t|t}$ and S_t as follows:

$$\begin{aligned} X_{t|t} &= AX_{t-1|t-1} + K(S_t - CAX_{t-1|t-1}), \\ S_t &= CAX_{t-1|t-1} + S_t - CAX_{t-1|t-1}, \end{aligned}$$

where the first equation follows from (20). Let Σ_S denote the variance-covariance matrix $Var_{t-1}[S_t]$ obtained from the Kalman filter. Suppose this matrix can be factorized as $\Sigma_S = GG'$ for some matrix G. Consider the model

$$\hat{X}_t = A\hat{X}_{t-1} + KG\hat{V}_t, \tag{24}$$

$$S_t = CA\hat{X}_{t-1} + G\hat{V}_t, \tag{25}$$

where $\hat{V}_t$ is an m-dimensional vector of mutually independent, i.i.d. standard normal shocks. Identifying $\hat{X}_t$ with $X_{t|t}$ and $\hat{V}_t$ with $S_t - CAX_{t-1|t-1}$ we obtain the following result.

Lemma 2 *For any matrix G that satisfies $GG' = \Sigma_S$ the original signal extraction model is observationally equivalent to (24)-(25) with the assumption that the agent perfectly observes the state $\hat{X}_t$ and the shock $\hat{V}_t$.*

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, 2004, "Exploring Pigou's Theory of Cycles," *Journal of Monetary Economics*, 51(6), 1183-1216.

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

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

[7] Blundell, R., Preston, I., 1998, "Consumption Inequality and Income Uncertainty," *Quarterly Journal of Economics*, 113(2): 603-640.

[8] 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.

[9] Christiano, L., C. Ilut, R. Motto and M. Rostagno, 2007, "Signals: Implications for Business Cycles and Monetary Policy," mimeo, Northwestern University.

[10] Cochrane, John H., 1994. "Shocks," *Carnegie-Rochester Conference Series on Public Policy*, Volume 41: 295-364.

[11] 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.

[12] Galí, J., 1992, "How Well Does the IS-LM Model Fit Postwar US Data?", *Quarterly Journal of Economics*, 107(2): 709-738.

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

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

[15] Jaimovich, N. and S. Rebelo, 2006, "Can News about the Future Drive the Business Cycle?" mimeo, Northwestern University.

[16] Leeper, E. M., T. B. Walker and S. S. Yang, 2009, "Fiscal Foresight and Information Flows," NBER Working Paper 14630.

[17] Lippi, F. and S. Neri, 2007, "Information variables for monetary policy in an estimated structural model of the euro area," *Journal of Monetary Economics* 54(4), 1256-70.

[18] Lorenzoni, G., 2009, "A Theory of Demand Shocks,"*American Economic Review*, 99(5), 2050-84.

[19] 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.

[20] Quah, D., 1991, "The Relative Importance of Permanent and Transitory Components: Identification and Some Theoretical Bounds," NBER Technical Working Papers 106.

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

[22] Schmitt-Grohé, S. and M. Uribe, 2008, "What's News in Business Cycles," mimeo, Columbia University.

[23] Shapiro, M.D. and M.W. Watson, 1988, "Sources of Business Cycle Fluctuations," *NBER Macroeconomics Annual*, MIT Press.

[24] Sims, C. and T. Zha, 2006, "Does Monetary Policy Generate Recessions?" *Macroeconomic Dynamics*, 10(2), 231-272.

[25] Sims, E., 2009, "Non-Invertibilities and Structural VARs," mimeo, Michigan.

6 Online appendix

6.1 More on non invertibility

This is the result discussed at the end of Section 2.2. The result is derived for the theoretical impulse responses, ignoring sampling error.

Proposition 1 *Suppose the econometrician observes (c_t, a_t, s_t) or (c_t, a_t). The impulse response of c_t to any identified shock from a structural VAR are either permanent and flat or zero.*

Proof. Let w_t be an identified shock, corresponding to a linear combination of current and past observables. Denote by $\mathcal{I}_t$ the consumer's information and by $\mathcal{I}_t^e$ the econometrician's information. Applying the law of iterated expectations we get

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

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

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

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

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

for all $k \geq 0$. ■

6.2 DSGE Model

6.2.1 Setup

Households. The household preferences are given by (13) in the text. The household budget constraint is

$$P_tC_t + P_tI_t + T_t + B_t + P_t\mathcal{C}(U_t)\bar{K}_{t-1} = R_{t-1}B_{t-1} + \Upsilon_t + \int_0^1 W_{jt}N_{jt}dj + R_t^kK_t \tag{26}$$

where P_t is the price level, T_t is a lump sum tax, B_t are holdings of one period bonds, R_t is the one period nominal interest rate, Υ_t are aggregate profits, W_{jt} is the wage of specialized labor of type j, R_t^k is the capital rental rate. The investment-specific technology parameter $d_t = \log D_t$ follows the stochastic process

$$d_t = \rho_d d_{t-1} + \varepsilon_{dt}.$$

ε_{dt} and all the variables denoted ε from now on are i.i.d. shocks.

Households choose consumption, bond holdings, capital utilization, and investment each period so as to maximize their expected utility subject to (26) and a standard no-Ponzi condition. Nominal bonds are in zero net supply, so market clearing in the bonds market requires $B_t = 0$.

Production and labor services. Consumption and investment are in terms of a final good which is produced by competitive final good producers using the CES production function

$$Y_t = \left(\int_0^1 Y_{jt}^{\frac{1}{1+\mu_{pt}}} dj \right)^{1+\mu_{pt}},$$

which employs a continuum of intermediate inputs. Y_{jt} is the quantity of input j employed and μ_{pt} captures a time-varying elasticity of substitution across goods, where $\log(1+\mu_{pt}) = \log\left(1+\mu_p\right) + m_{pt}$ and m_{pt} follows the process

$$m_{pt} = \rho_p m_{pt-1} + \varepsilon_{pt} - \psi_p \varepsilon_{pt-1}.$$

The production function for intermediate good j is

$$Y_{jt} = \left(K_{jt}\right)^{\alpha} \left(A_t L_{jt}\right)^{1-\alpha}, \tag{27}$$

where K_{jt} and L_{jt} are, respectively, capital and labor services employed. As in our baseline model, the technology parameter a_t follows the process (1)-(3) and the representative consumer does not observe x_t and z_t separately, but observes a_t and the signal s_t given by (6). However, here we treat explicitly the constant term in TFP growth by letting

$$A_t = \Gamma^t e^{a_t}.$$

Intermediate good prices are sticky with price adjustment as in Calvo (1983). Each period intermediate good firm j can freely set the nominal price P_{jt} with probability $1-\theta_p$ and with probability θ_p is forced to keep it equal to P_{jt-1}. These events are purely idiosyncratic, so θ_p is also the fraction of firms adjusting prices each period.

Labor services are supplied to intermediate good producers by competitive labor agencies that combines specialized labor of types in $[0,1]$ using the technology

$$N_t = \left[\int_0^1 N_{jt}^{\frac{1}{1+\mu_{wt}}} dj \right]^{1+\mu_{wt}},$$

where $\log(1+\mu_{wt}) = \log\left(1+\mu_w\right) + m_{wt}$ and m_{wt} follows the process

$$m_{wt} = \rho_w m_{wt-1} + \varepsilon_{wt} - \psi_w \varepsilon_{wt-1}.$$

Specialized labor wages are also sticky and set by the household. For each type of labor j, the household can freely set the price W_{jt} with probability $1-\theta_w$ and has to keep it equal to W_{jt-1} with probability θ_w.

Market clearing in the final good market requires

$$C_t + I_t + \mathcal{C}\left(U_t\right)\bar{K}_{t-1} + G_t = Y_t, \tag{28}$$

market clearing in the market for labor services requires $\int L_{jt}dj = N_t$.

Government spending and monetary policy. Government spending is set as a fraction of output and the ratio of government spending to output is $G_t/Y_t = \psi + g_t$, where g_t follows the stochastic process

$$g_t = \rho_g g_{t-1} + \varepsilon_{gt}.$$

Monetary policy follows the interest rate rule

$$r_t = \rho_r r_{t-1} + \left(1-\rho_r\right)\left(\gamma_\pi \pi_t + \gamma_y \hat{y}_t\right) + q_t \tag{29}$$

where $r_t = \log R_t - \log R$ and $\pi_t = \log P_t - \log P_{t-1} - \pi$, π is the inflation target, $\hat{y}_t$ is defined below and q_t follows the process

$$q_t = \rho_q q_{t-1} + \varepsilon_{qt}.$$

6.2.2 Optimality conditions

Households. Define the marginal utility of consumption

$$\Lambda_t = \frac{1}{C_t - hC_{t-1}} - \beta h E_t\left[\frac{1}{C_{t+1} - hC_t}\right]. \tag{30}$$

The consumers Euler equation is then

$$\Lambda_t = \beta R_t E_t\left[\Lambda_{t+1}\frac{P_t}{P_{t+1}}\right]. \tag{31}$$

The optimality conditions for $\bar{K}_{t-1}$ and I_t are

$$\Phi_t = \beta E_t\left[\Lambda_{t+1}\left(R^k_{t+1}U_{t+1} - P_t\mathcal{C}\left(U_{t+1}\right)\right)\right] + \left(1-\delta\right)\beta E_t\Phi_{t+1}, \tag{32}$$

$$P_t\Lambda_t = \Phi_t D_t\left[1 - \mathcal{G}_t - \frac{I_t}{I_{t-1}}\mathcal{G}'_t\right] + \beta E_t\left[\Phi_{t+1}D_{t+1}\left(\frac{I_{t+1}}{I_t}\right)^2\mathcal{G}'_{t+1}\right], \tag{33}$$

where Φ_t is the Lagrange multiplier on the capital accumulation constraint. The optimality condition for capacity utilization U_t is

$$R^k_t = \mathcal{C}'\left(U_t\right). \tag{34}$$

Firms. The optimality condition for the final good producer yields the demand for intermediate good j,

$$Y_{jt} = Y_t \left(P_{jt}/P_t\right)^{-\left(1+\mu_{pt}\right)/\mu_{pt}},$$

while the zero profit condition gives the final good price $P_t = (\int P_{jt}^{1/\mu_{pt}} dj)^{\mu_{pt}}$.

Cost minimization by the intermediate good producers and constant returns to scale imply that the nominal cost of producing Y_{jt} is $M_t Y_{jt}$, where M_t is the marginal cost

$$M_t = \alpha^{-\alpha}(1-\alpha)^{-(1-\alpha)} \left(R_t^k\right)^{\alpha} \left(W_t/A_t\right)^{1-\alpha}. \tag{35}$$

Cost minimization also implies that all intermediate good firms choose the same capital-labor ratio

$$\frac{K_t}{N_t} = \frac{\alpha}{1-\alpha}\frac{W_t}{R_t^k}. \tag{36}$$

Calvo pricing implies that the firms changing their price in period t choose the price P_t^* that maximizes

$$E_t\left[\sum_{\tau=0}^{\infty} \theta_p^{\tau}\beta^{\tau}\frac{\Lambda_{t+\tau}}{P_{t+\tau}}\left(P_t^* - M_{t+s}\right) Y_{jt+s}\right]$$

subject to $Y_{jt+\tau} = Y_{t+\tau}\left(P_t^*/P_{t+\tau}\right)^{-\left(1+\mu_{pt+\tau}\right)/\mu_{pt+\tau}}$. Their optimality condition is thus

$$E_t\left[\sum_{\tau=0}^{\infty} \theta_p^{\tau}\beta^{\tau}\frac{\Lambda_{t+\tau}}{P_{t+\tau}}\left(\frac{1}{\mu_{pt+\tau}} - \frac{1+\mu_{pt+\tau}}{\mu_{pt+\tau}}\frac{M_{t+\tau}}{P_t^*}\right) Y_{it+s}\right] = 0. \tag{37}$$

The problem of the labor agency is similar to that of the final good producer, so the demand for specialized labor of type j is

$$N_{jt} = N_t\left(W_{jt}/W_t\right)^{-(1+\mu_{wt})/\mu_{wt}},$$

and the price of labor services is $W_t = (\int W_{jt}^{1/\mu_{wt}} dj)^{\mu_{wt}}$. Calvo pricing for wage setters implies that the workers that can adjust their wage maximize

$$E_t\left[\sum_{\tau} \theta_w^{\tau}\beta^{\tau}\left[\frac{\Lambda_{t+\tau}}{P_{t+\tau}} W_t^* N_{jt+\tau} - \frac{1}{1+\zeta} N_{jt+\tau}^{1+\zeta}\right]\right],$$

and their optimality condition is

$$E_t\left[\sum_{\tau=0}^{\infty} \theta_w^{\tau}\beta^{\tau}\left(\frac{\Lambda_{t+\tau}}{P_{t+\tau}}\frac{1}{\mu_{wt+\tau}} - \frac{1+\mu_{wt+\tau}}{\mu_{wt+\tau}}\frac{N_{jt+\tau}^{\eta}}{W_t^*}\right) N_{jt+\tau}\right] = 0. \tag{38}$$

6.2.3 Log-linear approximation

First, we need to normalize some variables to ensure their stationarity. In particular, we define

$$\hat{c}_t = \log(C_t/A_t) - \log(C/A)$$

where C/A denotes the value of C_t/A_t in the deterministic version of the model in which A_t grows at the constant growth rate Γ. Analogous definitions apply to the quantities $\hat{y}_t, \hat{k}_t, \hat{\bar{k}}_t, \hat{i}_t$. The quantities N_t and U_t are already stationary, so we have

$$n_t = \log N_t - \log N,$$

and a similar expression for u_t. For nominal variables, we also need to take care of non-stationarity in the price level, so we define

$$\begin{aligned} \hat{w}_t &= \log\left(W_t/(A_t P_t)\right) - \log\left(W/\left(AP\right)\right), \qquad r_t^k = \log(R_t^k/P_t) - \log(R^k/P), \\ m_t &= \log(M_t/P_t) - \log(M/P), \qquad r_t = \log R_t - \log R, \qquad \pi_t = \log(P_t/P_{t-1}) - \pi. \end{aligned}$$

Finally, for the Lagrange multipliers we define

$$\hat{\lambda}_t = \log(\Lambda_t A_t) - \log\left(\Lambda A\right), \qquad \hat{\phi}_t = \log(\Phi_t A_t/P_t) - \log(\Phi A/P).$$

Notice that the hat is only used for variables normalized by A_t.

Optimality conditions. Conditions (30)-(36) can be log-linearized to yield the following seven conditions

$$\begin{aligned} \hat{\lambda}_t &= \frac{h\beta\Gamma}{(\Gamma - h\beta)(\Gamma - h)} E_t \hat{c}_{t+1} - \frac{\Gamma^2 + h^2\beta}{(\Gamma - h\beta)(\Gamma - h)} \hat{c}_t + \frac{h\Gamma}{(\Gamma - h\beta)(\Gamma - h)} \hat{c}_{t-1} + \\ &\quad + \frac{h\beta\Gamma}{(\Gamma - h\beta)(\Gamma - h)} E_t[\Delta a_{t+1}] - \frac{h\Gamma}{(\Gamma - h\beta)(\Gamma - h)} \Delta a_t, && (39) \\ \hat{\lambda}_t &= r_t + E_t\left[\hat{\lambda}_{t+1} - \Delta a_{t+1} - \pi_{t+1}\right], && (40) \\ \hat{\phi}_t &= (1-\delta)\beta\Gamma^{-1} E_t\left[\hat{\phi}_{t+1} - \Delta a_{t+1}\right] && (41) \\ &\quad + \left(1 - (1-\delta)\beta\Gamma^{-1}\right) E_t[\hat{\lambda}_{t+1} - \Delta a_{t+1} + r_{t+1}^k], && (42) \\ \hat{\lambda}_t &= \hat{\phi}_t + d_t - \chi\Gamma^2\left(\hat{i}_t - \hat{i}_{t-1} + \tilde{\Delta} a_t\right) + \beta\chi\Gamma^2 E_t\left[\hat{i}_{t+1} - \hat{i}_t + \tilde{\Delta} a_{t+1}\right], && (43) \\ r_t^k &= \xi u_t, && (44) \\ m_t &= \alpha r_t^k + (1-\alpha)\hat{w}_t, && (45) \\ r_t^k &= \hat{w}_t - \hat{k}_t + n_t. && (46) \end{aligned}$$

Resource constraints and inflation. Log-linearizing conditions (15) and (14) yields

$$\hat{k}_t = u_t + \hat{\bar{k}}_{t-1} - \Delta a_t, \tag{47}$$

$$\hat{\bar{k}}_t = (1-\delta)\,\Gamma^{-1}\left(\hat{\bar{k}}_{t-1} - \Delta a_t\right) + \left(1 - (1-\delta)\,\Gamma^{-1}\right)\left(d_t + \hat{i}_t\right). \tag{48}$$

Approximating and aggregating (27) over intermediate good producers and using the final good production function yields

$$\hat{y}_t = \alpha \hat{k}_t + (1-\alpha)\, n_t. \tag{49}$$

Market clearing in the final good market (28) yields

$$(1-\psi)\,\hat{y}_t = \frac{C}{Y}\hat{c}_t + \frac{I}{Y}\hat{i}_t + \frac{R^k K}{PY} u_t + g_t. \tag{50}$$

$C/Y, I/Y$ and $R^k K/(PY)$ are all equilibrium ratios in the deterministic version of the model in which A_t grows at the constant rate Γ.

Aggregating individual optimality conditions for price setters yields

$$\pi_t = \beta E_t \pi_{t+1} + \kappa m_t + \kappa m_{pt}, \tag{51}$$

where $\kappa = (1-\theta\beta)(1-\theta)/\theta$. Aggregating individual optimality conditions for wage setters yields

$$\begin{aligned} \hat{w}_t = {} & \frac{1}{1+\beta}\hat{w}_{t-1} + \frac{\beta}{1+\beta}E_t\hat{w}_{t+1} - \frac{1}{1+\beta}\left(\pi_t + \Delta a_t\right) + \frac{\beta}{1+\beta}E_t\left(\pi_{t+1} + \Delta a_{t+1}\right) \\ & -\kappa_w\left(\hat{w}_t - \zeta n_t + \hat{\lambda}_t\right) + \kappa_w m_{wt}, \end{aligned} \tag{52}$$

where

$$\kappa_w = \frac{(1-\theta_w\beta)(1-\theta_w)}{\theta_w(1+\beta)\left(1+\zeta\left(1+\frac{1}{\mu_w}\right)\right)}.$$

Summing up, we have 23 variables: the 9 exogenous variables

$$d_t, a_t, x_t, z_t, s_t, m_{pt}, m_{wt}, g_t, q_t,$$

and the 14 endogenous variables

$$\hat{c}_t, \hat{i}_t, \hat{y}_t, n_t, \hat{\bar{k}}_t, \hat{k}_t, u_t, \hat{\lambda}_t, \hat{\phi}_t, r_t^k, \hat{w}_t, \pi_t, m_t, r_t.$$

The model dynamics are given by the exogenous processes for the exogenous variables above and by the 14 equations (29) and (39)-(52).

The observables that are matched to the data are $\Delta y_t, \Delta c_t, \Delta i_t, \Delta w_t, \Delta n_t, r_t, \pi_t$, where the first four variables are obtained by adding $\Delta\hat{a}_t$ to the first differences of the corre-

sponding variables, e.g.,

$$\Delta y_t = \hat{y}_t - \hat{y}_{t-1} + \Delta \hat{a}_t.$$

6.2.4 New Keynesian model

The standard new Keynesian model is a special case of the DSGE model above when: (i) the capital stock and capacity utilization are fixed (and hence the investment-specific technology shock is absent), (ii) the habit parameter in the preferences h is set to zero, (iii) labor is homogeneous and wages are flexible. Furthermore, we consider a simple new Keynesian model with no fiscal policy ($G_t = 0$) and a simplified monetary rule with $\rho_r = \gamma_y = 0$ and no monetary shocks ($q_t = 0$).

In this case, the log-linearized model boils down to two stochastic difference equations which characterize the joint behavior of output and inflation in equilibrium:

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

where $y_t = \hat{y}_t + a_t$. The first comes from (29), (39), (40), (50). The second comes from (39), (45), (51) and the fact that (52) requires $\hat{w}_t = \zeta n_t - \hat{\lambda}_t$ when $\theta_w \to 0$. The coefficient $\tilde{\kappa}$ is given by $\left(1 + \zeta / \left(1 - \alpha\right)\right)\kappa$. As long as $\gamma_\pi > 1$ this system has a unique locally stable solution where y_t and π_t are linear functions of the four exogenous state variables $a_t, x_{t|t}, x_{t-1|t}, z_{t|t}$,

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

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

$$\begin{bmatrix} 1 & \gamma_\pi \\ -\tilde{\kappa} & 1 \end{bmatrix} D = \begin{bmatrix} 0 & 0 & 0 & 0 \\ -\tilde{\kappa} & 0 & 0 & 0 \end{bmatrix} + \begin{bmatrix} 1 & 1 \\ 0 & \beta \end{bmatrix} D \begin{bmatrix} 0 & 1+\rho_x & -\rho_x & \rho_z \\ 0 & 1+\rho_x & -\rho_x & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & \rho_z \end{bmatrix}.$$

The elements of D are a continuous non-linear function of $\tilde{\kappa}$ and a bit of algebra shows that

$$\lim_{\tilde{\kappa} \to 0} D = \frac{1}{1-\rho_x} \begin{bmatrix} 0 & 1 & -\rho_x & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}.$$

Since $\tilde{\kappa} \to 0$ when $\theta \to 1$, this shows that in the limit with very infrequent price adjustments we get (9).

6.3 Data

Our dataset spans the period 1954:3 to 2011:1. The start date is chosen because the effective Federal Funds Rate quarterly series, published by the Federal Reserve Board, is not available before that date.

The series for Real GDP, Real Personal Consumption Expenditures, Real Personal Durable Consumption Expenditures, Real Gross Private Domestic Investment, Wages and the GDP Implicit Price Deflator are from the Bureau of Economic Analysis (available through the Federal Reserve Bank of Saint Louis online database). Population and employment series are from the Bureau of Labor Statistics online database (series IDs LNS10000000Q and LNS12000000Q respectively). The Federal Funds Rate series is from the Federal Reserve Board online database (series ID H15/H15/RIFSPFF N.M).

The GDP series is constructed by dividing Real GDP by population. The consumption series is constructed by subtracting Real Personal Durable Consumption from Real Personal Consumption and dividing by population. The investment series is constructed by dividing the sum of Real Gross Investment and Real Personal Durable Consumption by population. The labor input series is constructed by dividing Employment by Population. Inflation is constructed by computing the quarterly log difference of the Price Deflator. The real wage is constructed by dividing Real Wages by the Price Deflator. The nominal interest rate is the effective Federal Funds Rate.

6.4 More on recovering states and shocks

This section presents more results on the Kalman smoother for the baseline model, following up on Section 2.5. Figure 8 plots the RMSE of the smoothed estimates of x_t and z_t, when data up to $t+j$ are available, for $j = 0, 1, 2, \ldots$. The RMSE is 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 the RMSE at the steady state of the Kalman filter, that is, assuming the forecaster has access to 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 ob-

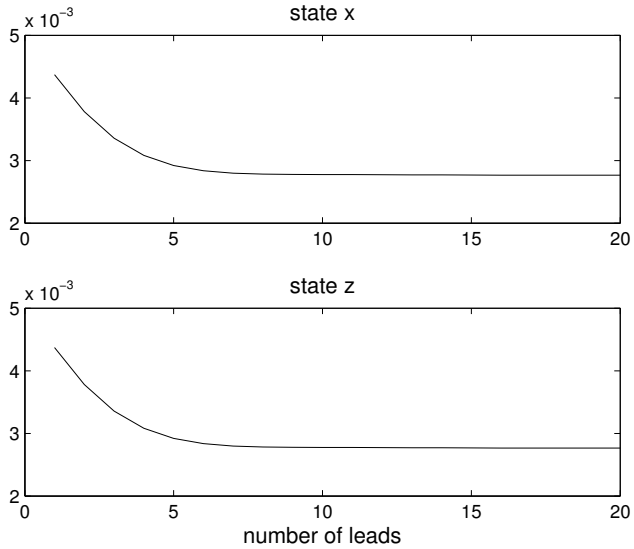

Figure 8: RMSE of the estimated states at time t using data up to $t + j$

servations 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 is 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 8 illustrates the results discussed in the text. Notice that most of the relevant information arrives in the first six quarters after t there are minimal gains after that.

Figure 9 we report the RMSE of the estimates of the shocks ϵ,η and ν. To help the interpretation, each RMSE is normalized dividing it by the ex ante standard deviation of the respective shock (e.g., by σ_ϵ for ϵ).

Notice that if the model was invertible, all RMSE would be zero at $j = 0$. The fact that all RMSE are bounded away 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 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

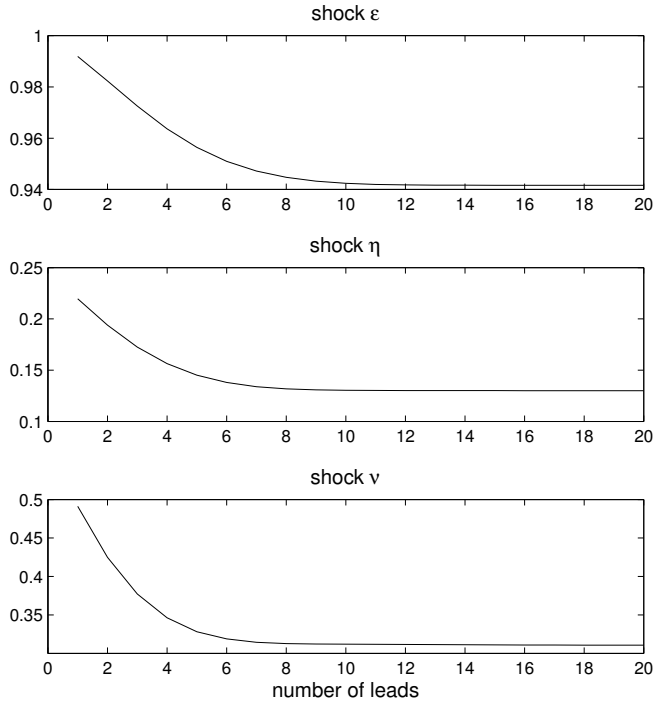

Figure 9: Normalized RMSE of the estimated shocks at time t using data up to $t + j$

relatively precise estimates of the state x_t, as seen in Figure 8? 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.