Essays on Imperfect Information, Macroeconomic Fluctuations, and Nominal Rigidities

by

Jean-Paul L'Huillier

B.A., University of Geneva (2002)

Submitted to the Department of Economics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Economics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2010

Author ....
Department of Economics
May 15, 2010

Certified by ....
Olivier J. Blanchard
Class of 1941 Professor of Economics
Thesis Supervisor

Accepted by ....
Esther Duflo
Abdul Latif Jameel Professor of Poverty Alleviation
and Development Economics
Chairman, Departmental Committee on Graduate Theses

Essays on Imperfect Information, Macroeconomic Fluctuations, and Nominal Rigidities

by

Jean-Paul L'Huillier

Submitted to the Department of Economics
on May 15, 2010, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Economics

Abstract

The first essay 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.

The second essay experimentally examines whether looking at other people's pricing decisions is a type of heuristic, a decision rule that people over-apply even when it is not applicable. such as in the case of clearly private value goods. We find evidence that this is indeed the case. individual valuation of a purely subjective experience under full information, elicited using incentive compatible mechanism, is highly influenced by values of others. As the third essay shows, this result can shed light on price rigidities.

Inspired by the experimental results of the second essay, the third essay develops a model of slow macroeconomic adjustment to monetary shocks. The model exploits the idea that buyers are imperfectly informed about their nominal valuation. I proceed in three steps. First, I develop a mechanism for price rigidities. My mechanism captures the notion that firms are reluctant to increase prices after an increase in demand or costs because it creates a disproportionate adverse reaction among consumers. These reactions arise endogenously for purely informational reasons. The key assumption is that some consumers are better informed than others about monetary shocks. If few consumers are informed, equilibria with nominal rigidity exist. In these equilibria

firms do not change prices even though they are arbitrarily well informed, and have no menu costs. Moreover, if the proportion of informed consumers is low enough, these equilibria dominate equilibria with flexible prices. Second, I show that when firms do not change prices they inflict an informational externality on other firms. Consumers buy goods sequentially, one after the other, and change their beliefs about shocks when they see prices change. Therefore, when firms do not change prices, consumers do not learn. This hurts both firms and consumers. Third, I study the dynamic responses of output and inflation to shocks. Because of the informational externality learning is initially slow, the responses are delayed and hump-shaped. The responses are also asymmetric – prices increase faster than they decrease, and therefore negative shocks trigger larger output responses than positive shocks.

Thesis Supervisor: Olivier J. Blanchard
Title: Class of 1941 Professor of Economics

To my mother, who at very early stages of my life provided me with infinite love, and thus gave me the sense of security to go where I have gone.

Acknowledgments

This thesis is the achievement of an incredible intellectual journey that allowed me to link seemingly unrelated findings and methodologies under a common theme: imperfect information. In such a journey it is easy to get lost, and I am most thankful to my three advisors for having so skillfully guided my thoughts.

I have been lucky enough to have been guided by a very diverse committee. In an effort to be specific, I benefited from Olivier Blanchard's powerful intuition to detect good ideas, and I thank him for highlighting that not all ideas are worth working on. I thank Mikhail Golosov for teaching me how to be pragmatic and transform an idea into a modern full-fleshed macroeconomic model, and then into a paper. I thank Guido Lorenzoni for showing me the benefits of patient and careful thinking.

Being a student at the MIT Economics Department is a truly fantastic experience. I have benefited enormously from numerous exchanges with my classmates, other students, and many faculty members. I cannot possibly say enough to thank all of them.

Several people had a particularly crucial impact leading to the success of this thesis. Among them, I would like to thank Pablo Kurlat, Matthew Notowidigdo, Alain Quiamzade, Roberto Rigobon, Paul Schrimpf, Birger Wernerfelt, and Muhamet Yildiz for their support and brilliant insights.

News, Noise, and Fluctuations: An Empirical Exploration

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

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

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

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:

$$\rho_x = \rho_z = \rho,$$

$$\sigma_\epsilon^2 = (1-\rho)^2 \sigma_u^2, \qquad \sigma_\eta^2 = \rho \sigma_u^2,$$

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

[1] A 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.

[2] To 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

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

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

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

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

Combining the last two equations gives

$$c_t = \lim_{j\to\infty} E_t[a_{t+j}]. \tag{6}$$

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

$$s_t = x_t + \nu_t, \tag{7}$$

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

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

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:

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

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 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.

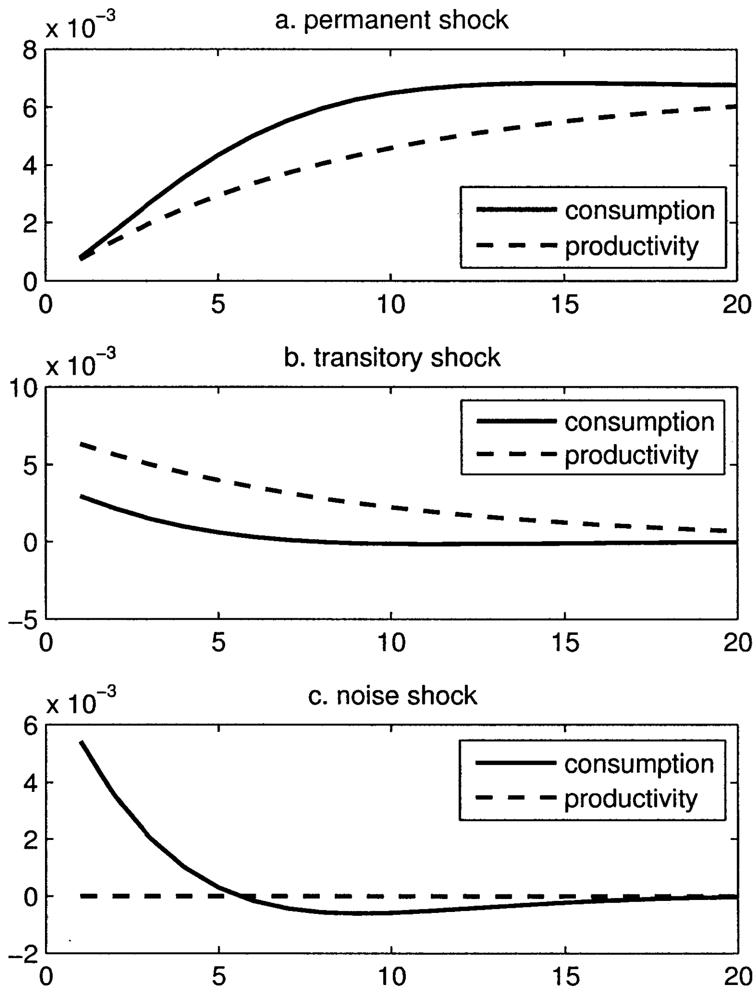

Figure 1: Impulse Responses to the Three Shocks

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

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

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

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

with

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

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

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

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

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

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

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

$$E_{t-1}[u_t^c] = \frac{1}{1-\rho} E_{t-1}[\epsilon_t] = 0,$$

$$E_{t-1}[u_t^a] = E_{t-1}[\epsilon_t + \eta_t] = 0.$$

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 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 at least c_t and a_t, past productivity does not help predict consumption, but past consumption typically helps to predict productivity as it captures the consumers' information on the permanent component x_t.[4]

[4]The 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, we 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 combination 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 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 $c_t = \lim_{j\to\infty} E\left[a_{t+j}|\mathcal{I}_t\right]$.

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

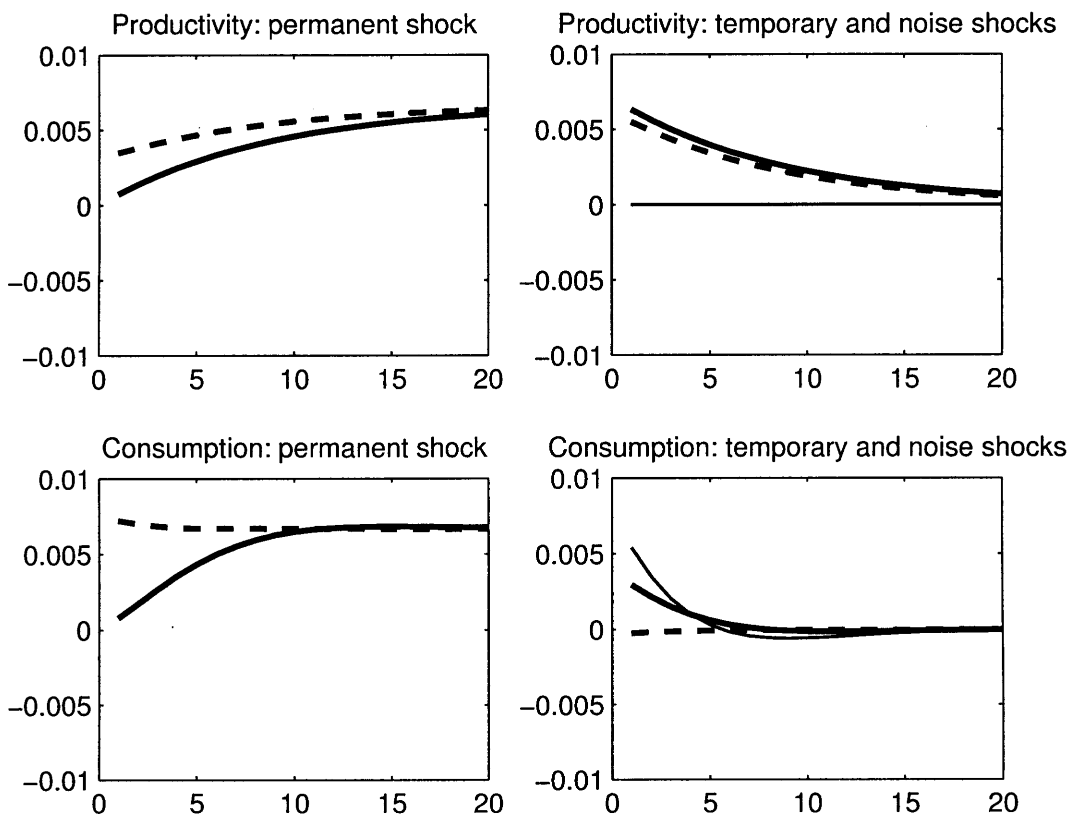

Figure 2: True and SVAR-based estimated impulse responses

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

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

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 that the correlation is higher the higher the 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

[6]The 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.

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

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

[8]This 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

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

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 efficiently all the information in the sample 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 a different moment: the coefficient of correlation 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.

[9]These 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.

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]

- 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

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

[11] For 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

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 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;

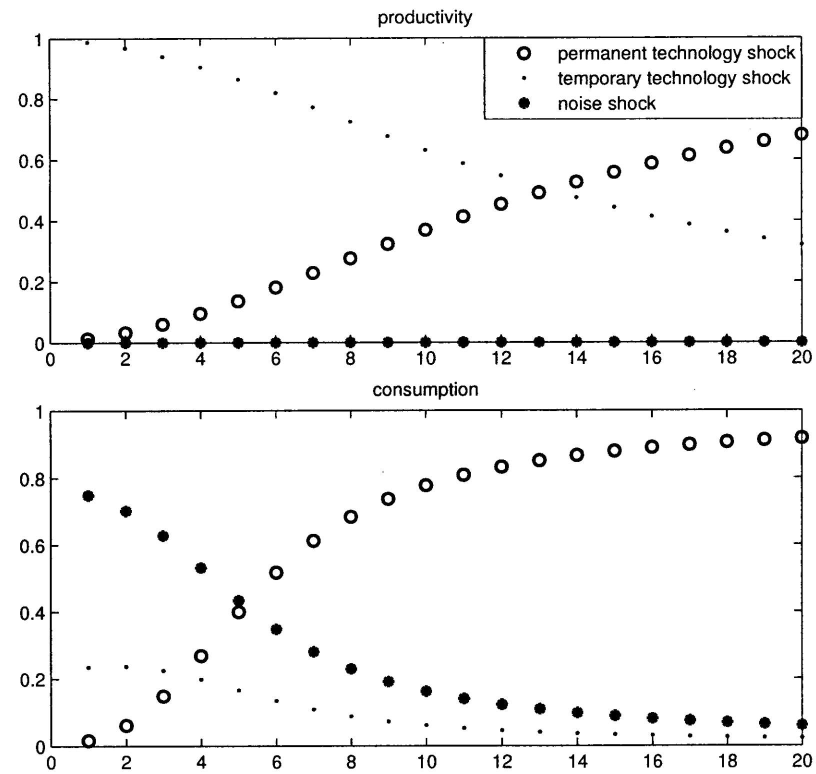

Figure 3: Variance Decomposition: Benchmark Model

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, 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 from 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 "smoothed" 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 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. Some times 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 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 $j = 0, 1, 2, ...$ 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. This is a numerical result: computing the RMSE of s_t from the econometrician's Kalman smoother, we find that it is equal 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%.

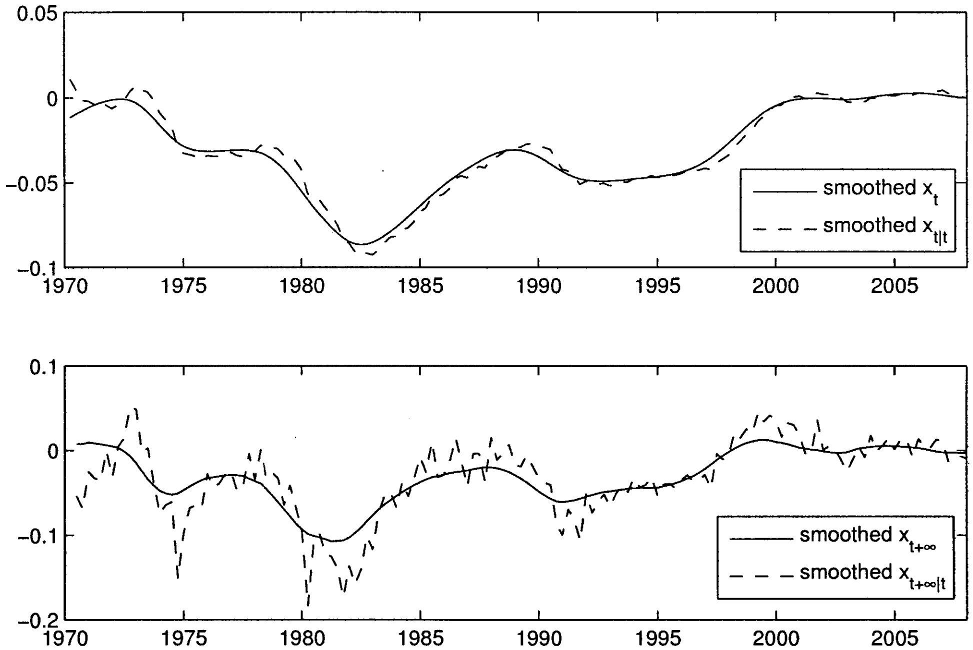

Figure 4: Smoothed estimates of the permanent component of productivity, of long-run productivity, and of consumers' real time expectations
Top panel: smoothed estimate of x_t (solid line) and of $x_{t|t}$ (dashed line)
Bottom panel: smoothed estimate of $x_{t+\infty}$ (solid line) and of $x_{t+\infty|t}$ (dashed line)

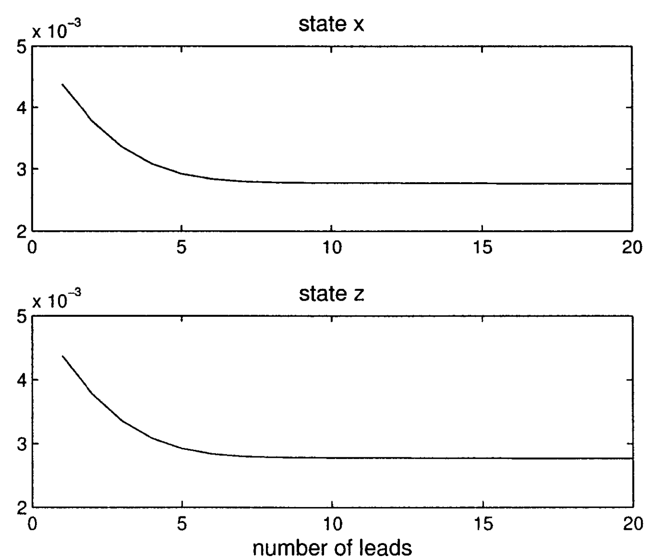

Figure 5: RMSE of the estimated states at time t using data up to $t + j$

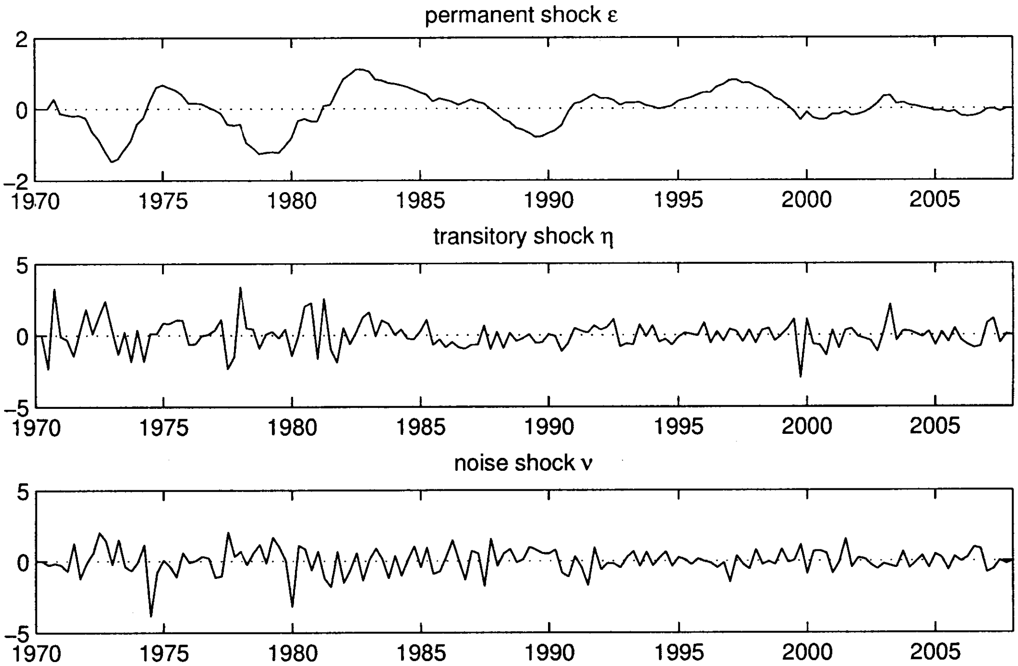

Figure 6: Smoothed estimates of the shocks

However, most of the relevant information arrives in the first six quarters, after that, there are minimal gains in the precision of the estimate.

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.

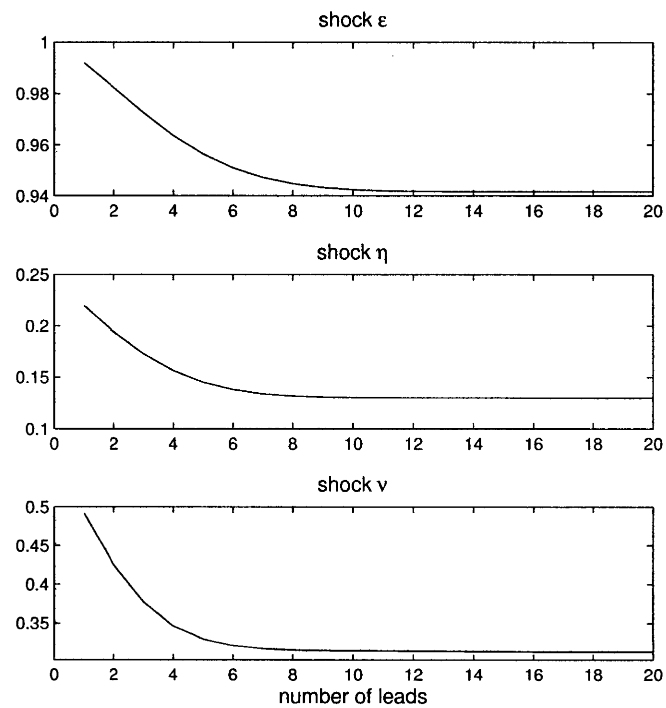

Figure 7: Normalized RMSE of the estimated shocks at time t using data up to $t + j$

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 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 processes for the permanent and transitory components of productivity, the information structure, the behavior of consumers. And, unfortunately, the estimated parameters are likely to be sensitive to the specific assumptions.

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

The first is motivated by the data. As we saw from Table 1, the dynamics of consumption and the dynamic relation between productivity and consumption are richer than those implied by the benchmark. These require at least a modification of our assumptions about consumption behavior. Our assumption about consumption implies that consumption 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

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

where the coefficients d are non-linear functions of the following model 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

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

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

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.

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):

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

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).

δ	ρ	σ_u	σ_ϵ	σ_ζ	σ_ν	ML
0	0.8785	0.0068	0.0008	0.0063	0.0086	1073.3
0.1	0.8700	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_ν	ML
0	0.8910	0.0067	0.0089	1073.2
0.1	0.8989	0.0069	0.0067	1072.9
0.2	0.9110	0.0072	0.0052	1071.6
0.3	0.9249	0.0077	0.0039	1068.2
0.4	0.8948	0.0085	0.0	1064.6
0.6	0.9434	0.0114	0.0	1034.1
0.8	0.9645	0.0229	0.0	937.6
1	0.0070	0.0067	0.0857	391.1

Table 7: Maximum Likelihood Estimation: Labor hoarding

Our second extension is motivated by the discussion of labor hoarding and pro-cyclical productivity in the research on the relation between output and employment. Our benchmark model has assumed that labor productivity is exogenous; there is however substantial evidence 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 6 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 compensate 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^c and u^a. Allowing for endogenous productivity, gives us an alternative channel to explain this correlation. The results in Table 6 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 lung 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, as 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

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

with

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

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

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

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

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

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 $a_t, x_{t|t}, x_{t-1|t}, z_{t|t}$,

$$\left(\begin{array}{c} y_t \\ \pi_t \end{array}\right) = D_\kappa \left(\begin{array}{c} a_t \\ x_{t|t} \\ x_{t-1|t} \\ z_{t|t} \end{array}\right).$$

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

$$\left[\begin{array}{cc} 1 & \phi \\ -\kappa & 1 \end{array}\right] D_\kappa = \left[\begin{array}{cccc} 0 & 0 & 0 & 0 \\ -\kappa & 0 & 0 & 0 \end{array}\right] + \left[\begin{array}{cc} 1 & 1 \\ 0 & \beta \end{array}\right] D_\kappa \left[\begin{array}{cccc} 0 & 1+\rho & -\rho & \rho \\ 0 & 1+\rho & -\rho & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & \rho \end{array}\right].$$

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

$$\lim_{\kappa \to 0} D_\kappa = \frac{1}{1-\rho} \left[\begin{array}{cccc} 0 & 1 & -\rho & 0 \\ 0 & 0 & 0 & 0 \end{array}\right].$$

Since $\kappa \longrightarrow 0$ when $\theta \longrightarrow 1$, 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 $P \equiv Var_{t-1}[\xi_t]$. 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

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

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 $k \geq 0$.

Appendix D. Econometrician's Kalman Filter

The econometrician's state vector is given by

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

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

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

Then the state ξ_t^E evolves according to:

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

where the matrices Q and R are given by

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

$$R = \begin{bmatrix} & 1 \ 0 \ 0 & \\ & 0 \ 0 \ 0 & \\ & 0 \ 1 \ 0 & \\ B\begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix} & +B\begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix} & +B\begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} \end{bmatrix}.$$

When the econometrician can observe (a_t, c_t), the observation equation is, in matrix form,

$$(a_t, c_t)' = TX_t, \tag{20}$$

where

$$T = \begin{bmatrix} 1 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 + \frac{\rho_x}{1-\rho_x} & -\frac{\rho_x}{1-\rho_x} & 0 \end{bmatrix}.$$

The econometrician's filtering problem can then be solved from (19)-(20). The case in which the econometrician can also observe s_t is treated in a similar way. This filter can be used both to compute recursively the likelihood function and to derive smoothed estimates of the unobservable states in ξ_t^E, as in Section 4.4. Expanding the state space to include the shocks $(\epsilon_t, \eta_t, \nu_t)$, it is easy to compute their smoothed estimates, as in Section 4.5.

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.

Public and Private Values

Dan Ariely, Anat Bracha, and Jean-Paul L'Huillier*

Abstract

In this paper we experimentally examine whether looking at other people's pricing decisions is a type of heuristic, a decision rule that people over-apply even when it is not applicable . such as in the case of clearly private value goods. We find evidence that this is indeed the case . individual valuation of a purely subjective experience under full information, elicited using incentive compatible mechanism, is highly influenced by values of others. This result can shed light on price behavior, price rigidities and rents.

Keywords: second price auction, heuristic.

*Duke University, Tel Aviv University, and MIT. We thank Peter Diamond and Andrei Shleifer for useful comments, and Daniel Mochon and Ray Weaver for help during the experiments.

Introduction

In many real-life situations before we make our own decisions it is useful to look at what others have done when they were facing similar situations. For instance, when considering a job offer we might consult others who recently accepted or rejected a related job offer. When purchasing an apartment we might seek information about what prices were paid recently for similar apartments, and if possible in the same vicinity or even the same building.

This general strategy is often a very useful one, as many, if not most, of the situations we encounter include some component of common values. This means that relying on others' decisions, or the prices they are willing to pay, can be helpful for our decisions, as they convey some relevant information. However, although this strategy is often a good one, at times applying it can be inappropriate. One notable example where following others is inappropriate is information cascades, where individuals fail to realize the information externality of their actions (Banerjee, 1992; Bikhchandani, Hirshleifer, and Welch, 1992; Simonsohn and Ariely, 2008). Another is Asch's famous "line" experiment (Asch 1951) demonstrating that even in visual judgments the majority can influence individuals, which is common in a range of social behaviors (See Petty and Wegener 1998 for a survey).

In this paper we examine if relying on the behaviors of others can also occur in the case of pricing decisions for goods that are private values goods. Specifically, we ask whether looking at other people's pricing decisions is a type of heuristic, a decision rule that people over-apply even when this information is not applicable -- such as in the case of clearly private value goods. We provide experimental evidence showing that bids in private-value auctions indeed depend on the bids provided by others -- a relationship we expect to find in common-value but not in private-value auctions.

To ensure purely private value situations and full information we use a novel hedonic stimulus -- annoying sounds. This type of experience provides us with an experience that is direct, hedonic, providing participants with full information, and importantly it is the type of experience that has no externalities and no possible trade. Elicitation of values in our experiment is done using an incentive-compatible mechanism, either BDM procedure (in stage 1 of the experiment) or a second- price auction (in stage 2 of the experiment). We find that subjective values, or bids, are interrelated in the sense that observing others' past bids had a statistically significant effect on the bids participants submitted.

One implication of these results is that figuring out how much we are willing to pay, even for simple experiences, may be more complex, or less readily available, than we usually consider. As a consequence of this complexity we over-apply the strategy of using others' behavior and use it as an input for our decisions. This tendency may explain some anomalous price behavior and fashion trends.

The Experiment

To make the point that participant's values are interrelated even in situation of private values with perfect information, no externalities and with no trade, we recruited 45 MIT students. There were two stages to the experiment, with ten rounds each; in both stages, and in all rounds, participants bid for listening to an unpleasant tone over a headset. We used this artificial hedonic "good" for the following reasons: (1) we were able to provide subjects with a sample of the tone before they subsequently placed their bid, and as a result create a perfect information situation, (2) this is a purely subjective and personal experience, (3) there are no externalities, (4) there are no trading possibilities, and (5) to avoid a situation in which subjects could solve the pricing problem drawing on their own experiences. Unpleasant tones are not traded in the marketplace, so our subjects could not refer to similar decisions made outside the laboratory as a basis for their valuation. These aspects of the stimuli enable us to focus on the effect we are interested in, if exists (the procedure is similar to Ariely et al (2003)).

In the experiment we had two stages: in stage 1 we recorded actual bids of a first group of participants. These bids where then used in stage 2 to study their influence on a second group of participants.

Stage 1

Upon arrival to the lab, at the beginning of stage 1, we explained that the study involves bidding and listening to high-pitched sound. We then explained the bidding procedure – the BDM (Becker et al 1964) procedure – a second price auction against a random number. In our particular case the random number was drawn for a uniform distribution between 5 cents and 95 cents. According to this procedure, if a subject's bid was lower than the number drawn, the subject would win the auction – he would listen to the sound and get paid the random number drawn. Otherwise, he would not listen to the sound, and would not get paid. This procedure was explained in detail before the start of the experiment, and we stressed its incentive compatibility property, including a few examples to show that bidding one's true value is the optimal strategy. Then, subjects listened to the sound on their headphones for 10 seconds, and immediately after this first experience with the tone, participants were asked to bid. That is to state the minimum amount of money they demand in order to listen to the same tone again (Willingness to Accept or WTA). This pricing procedure was repeated nine times (ten trials).

Table 1 below presents summary statistics of bids on each trial.

[Table 1 here]

Selection

Out of the 22 participants in stage 1, 17 bidders expressed consent to use their bids in future auctions; of these 17 participants we retained 12, avoiding bidders who consistently bid extreme values such as zero. The summary statistics of the bids of the 12 selected bidders are shown below.

[Table 2 here]

We then sorted the selected bids in each trial from low value to high value, and used the four lowest and four highest bids in each trial for the LOW and HIGH treatment of stage 2, as is explained next.

Stage 2

Stage 2 was similar to stage 1. That is, participants are asked to bid for listening to the same unpleasant high-pitched tone on a headset as in stage 1. Again, participants were first exposed to the tone, and only then asked to place a bid. However, in stage 2 subjects participated in a second price auction, instead of the BDM procedure used in Stage 1. In the instructions, the notion of second price auction was explained, and it was clarified that it is an incentive compatible mechanism, meaning that their best strategy is to state their true value. Subjects were told they are about to participate in several trials of second price auction against four other bidders, and that these bidders already placed their bids a few days earlier. The instructions followed: "In order to become familiar with the task, bids in the first round will be displayed to you momentarily...you will start bidding only in the second round."

Hence, starting in the second round subjects were asked to place their bid, i.e., to state the minimum amount of money they demand in order to listen to the same tone again (WTA). If the subject won the auction, she listened to the tone and earned the second-lowest bid in the auction (equal or greater than her bid). After each auction, the other four bids were shown to the subject. This procedure was repeated nine times, to complete ten trials.

We had two treatments -- HIGH and LOW: in each trial subjects participated in a second price auction against either the four highest (HIGH treatment) or four lowest (LOW treatment) bids of retained bids for the same trial in stage 1. For example, in the first trial the participating bids in the HIGH treatment were [100, 99, 60, 80], and in the LOW treatment were [20, 25, 5, 20]. In subsequent trials (trial 2 through 10) the set of bids retained from stage 1 in each treatment (HIGH and LOW) was very similar to the bids shown above.

At the end of stage 2 we asked for feedback, asking the subjects what was their bidding strategy. Out of the 23 participants who participated part in stage 2, one indicated s/he did not understand the auction and was confused. Another indicated s/he was testing the program to see whether the study was for real by, for instance, bidding 500 in one trial. As a result, we exclude these two participants when analyzing the data.

The results of Stage 2 are shown in Figure 1 with mean bids in HIGH and LOW treatments, and Table 3 showing these differences are significant in all trials ($p<.03$).

[Figure 1 Here]

[Table 3 Here]

Discussion

Our question is whether looking at other people's pricing decisions is a type of heuristic, a decision rule that people over-apply even when it is not applicable – such as in the case of clearly private value goods. We find evidence that this is indeed the case – individual valuation of a purely subjective experience is highly influenced by others' value.

These results might be useful when thinking about value formation and pricing in actual markets. If valuations of others affect individual's valuation this can lead to rigidities (of price or quantities), and create seemingly common value even in the absence of one, as in the case of fashion. This adds an aspect to firm-consumer interaction and gives an additional explanation as for why firms hire experts or public-opinion shapers. Beyond the role of information: firms may hire experts or public-opinion shapers to *generate* consumers' value and, in turn, economic rents.

This view of interrelated private valuation not only has interesting Industrial Organization and Marketing implications, as described above, but also Macroeconomics ones. For instance, interrelated private valuation can help explain the origins of nominal rigidities: consumers who cannot translate their valuation into monetary units, as is the case in our experiment, may rely on prices others are willing to pay. In such a case prices will be sticky.

Reference

Ariely, D, G. Loewenstein, and D. Prelec, 2003, Coherent Arbitrariness: Stable Demand Curves Without Stable Preferences, The Quarterly Journal of Economics 118, 73-105.

Asch, S. E., 1951, Effects of Group Pressure Upon the Modification and Distortion of Judgment, in Guerzkow, Ed., Groups, leadership and men (Carnegie Press, Pittsburgh) 177-190.

Banerjee, A. V., 1992, A Simple-Model of Herd Behavior, The Quarterly Journal of Economics 107, 797-817.

Becker, G. M., M. H. DeGroot, and J. Marschak, 1964, Measuring Utility by a Single-Response Sequential Method, Behavioral Science 9, 222-232.

Bikhchandani, S., D. Hirshleifer, and I. Welch, 1992, A Theory of Fads, Fashion, Custom, and Cultural-Change as Informational Cascades, Journal of Political Economy 100, 992- 1026.

Petty, R. E. and D. T.Wegener, 1998, Attitude change: Multiple roles for persuasion variables, in D. Gilbert, S. Fiske, and G. Lindzey, eds., The handbook of social psychology, 4th Edition (McGraw Hill, Boston) 323-390.

Simonsohn, U., and D. Ariely, 2008, When Rational Sellers Face Nonrational Buyers: Evidence from Herding on eBay, Management Science 54, 1624–1637.

Table 1: Summary Statistics

	Average	**S.d.**	**Min**	**Max**
Trial 1	41	37	0	100
Trial 2	37	34	0	100
Trial 3	36	33	0	100
Trial 4	32	28	0	100
Trial 5	35	33	0	100
Trial 6	31	29	0	100
Trial 7	35	32	0	100
Trial 8	34	31	0	100
Trial 9	34	30	0	100
Trial 10	33	32	0	100
Overall	**35**	**32**	**0**	**100**

Table 2: Selection

	Average	**S.d.**	**Min**	**Max**
Trial 1	49	32	5	100
Trial 2	43	27	15	100
Trial 3	42	27	15	95
Trial 4	39	22	10	80
Trial 5	41	28	10	95
Trial 6	38	26	10	80
Trial 7	42	27	5	80
Trial 8	42	26	5	77
Trial 9	40	24	5	77
Trial 10	43	30	1	90
Overall	**42**	**26**	**1**	**100**

Table 3: Main Results

	HIGH	LOW	p-val
Trial 2	37.3	11.7	0.022
Trial 3	30.5	9.2	0.015
Trial 4	28.9	11.3	0.027
Trial 5	27.4	9.9	0.015
Trial 6	28.1	8.7	0.009
Trial 7	26.6	8.4	0.011
Trial 8	27.1	10.6	0.025
Trial 9	28.5	8.9	0.010
Trial 10	26.5	8.2	0.019
Overall	28.97	9.65	

Figure 1: Main Results

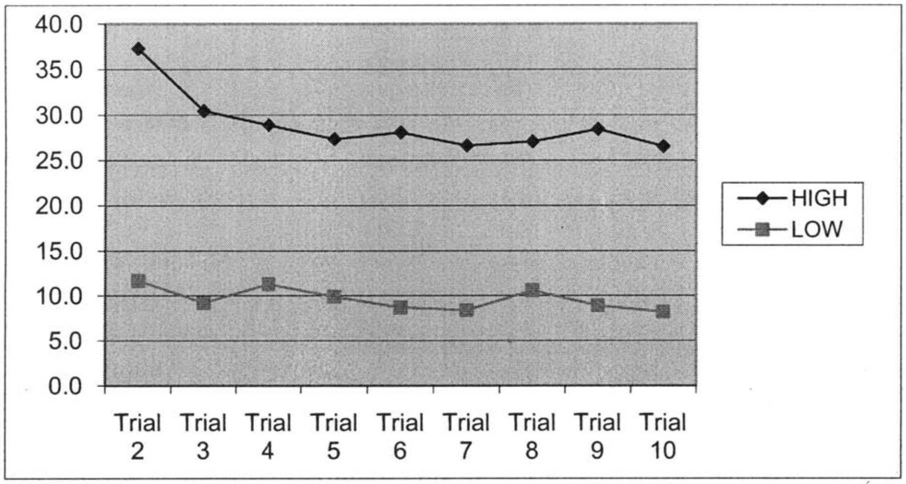

Consumers' Imperfect Information and Nominal Rigidities

Jean-Paul L'Huillier*

January 2010

Abstract

This paper develops a model of slow macroeconomic adjustment to monetary shocks. I proceed in three steps. First, I develop a mechanism for price rigidities. My mechanism captures the notion that firms are reluctant to increase prices after an increase in demand or costs because it creates a disproportionate adverse reaction among consumers. These reactions arise endogenously for purely informational reasons. The key assumption is that some consumers are better informed than others about monetary shocks. If few consumers are informed, equilibria with nominal rigidity exist. In these equilibria firms do not change prices even though they are arbitrarily well informed, and have no menu costs. Moreover, if the proportion of informed consumers is low enough, these equilibria dominate equilibria with flexible prices. Second, I show that when firms do not change prices they inflict an informational externality on other firms. Consumers buy goods sequentially, one after the other, and change their beliefs about shocks when they see prices change. Therefore, when firms do not change prices, consumers do not learn. This hurts both firms and consumers. Third, I study the dynamic responses of output and inflation to shocks. Because of the informational externality learning is initially slow, the responses are delayed and hump-shaped. The responses are also asymmetric – prices increase faster than they decrease, and therefore negative shocks trigger larger output responses than positive shocks.

Keywords: signaling, distortion, decentralized trading, islands, dispersed information.

*MIT. I am highly indebted to my advisors Olivier Blanchard, Michael Golosov, and Guido Lorenzoni for generously providing me with guidance and suggestions that significantly improved the paper. I also thank Daron Acemoglu, Marios Angeletos, Nick Bloom, Ricardo Caballero, Arnaud Costinot, Giovanni Dell'Ariccia, Tatyana Deryugina, Mathias Dewatripont, Peter Eso, Jonathan Goldberg, Veronica Guerrieri, Nathan Hendren, Oleg Itskhoki, Larry Jones, Greg Kaplan, Anton Kolotilin, Pablo Kurlat, Plamen Nenov, Matt Notowidigdo, Mike Powell, Roberto Rigobón, David Romer, Aleh Tsyvinski, Kenichi Ueda, Iván Werning, Pierre Yared, and Muhamet Yildiz for very useful comments and suggestions.

1 Introduction

There is a long standing idea in economics that, when demand or costs increase, firms are reluctant to increase prices because it triggers a disproportionate adverse reaction among consumers. According to some authors, this idea could lie at the root of one of the main puzzles in macroeconomics, namely the existence of nominal price rigidities and real effects of money[1]. Interestingly, when asked to explain their reluctance to increase prices after an increase in costs, firms' managers usually answer that "price increases cause difficulties with customers" (Blinder et al. 1998). However, the reasons remain so far unclear. One possible explanation is that, for a variety of reasons, consumers question whether price increases are justified (Rotemberg 2005) or exhibit a strong distaste for them (Nakamura and Steinsson 2005; Heidhues and Köszegi 2005). One important issue left open by this solution, when considering price increases following an increase in money, is that rational consumers should understand that these increases are only nominal, not real. My paper addresses this issue and provides a model in which firms' difficulty to increase nominal prices arises for purely informational reasons.

I consider a dynamic economy that is subject to exogenous aggregate shocks. The key assumption is that there is heterogeneity of information regarding these shocks among consumers, some consumers being better informed than others. In my model, firms are able to learn the realization of shocks from the environment. By letting firms to have arbitrarily precise information, and without menu costs, I ask whether they adjust prices. My general goal is to study how information is transmitted in this economy, and its implications for output and prices. For concreteness, in this paper I focus on the case of monetary shocks, i.e. changes in the supply of money, but the basic idea can also be applied to other cases.

The environment is as follows. I consider an economy populated by firms and consumers. In this economy, goods markets are decentralized. More specifically, the economy is composed of islands, and on each there is a price-setting firm. Consumers travel from one island to the other, buying goods sequentially. Importantly, every period, the only price observable to consumers is the price of the island. After a finite number of periods, all consumers buy a good sold in a centralized and competitive economy-wide market. In a cash in advance environment, the price of this good is determined by the supply of money. Thus, this centralized market is meant to capture the idea that in the long run prices are flexible, and do not reflect any strategic concerns between suppliers and buyers.

After a monetary shock, an exogenous proportion of consumers learns its realization. At this point, all other consumers and firms are uninformed. Firms, being so far uninformed, keep

[1]From the literature it is possible to see that this idea goes back to at least Hall and Hitch (1939), and has been mentioned by many other authors. For some examples, see Okun (1981); Kahneman, Knetsch, and Thaler (1986); Greenwald and Stiglitz (1989); Fehr, Kirchsteiger, and Riedl (1993); Anderson and Simester (forthcoming).

their price unchanged but sell as much consumers demand. Because every island is visited by a representative sample of consumers, firms are able to learn the realization of the shock by observing total demand. This gives rise to an informational asymmetry between firms and uninformed consumers.

I obtain three sets of results. First, the informational asymmetry between firms and uninformed consumers gives rise to an interesting strategic tension. Uninformed consumers revise their beliefs about money shocks as a function of firms' prices. Firms make higher profits when consumers believe that monetary shocks are high, because they expect to face higher prices in the long run, and therefore are willing to spend more in the short run. These two facts lead to a *non existence* of equilibria mimicking perfect information when the proportion of informed is low. For the intuition, suppose that all consumers are uninformed, and that such an equilibrium exists. In this alleged equilibrium, prices are proportional to money supply – higher after positive money shocks than after negative shocks, and all consumers update their beliefs from prices. Thus, consumers spend more after positive shocks, and firms make higher nominal profits. But this means that when shocks are negative firms have a profitable deviation – to post high prices – and therefore this is not an equilibrium. Moreover, as it will become clear, the strategic tension between firms and uninformed consumers changes the nature of price increases, in particular, if only few consumers are informed, it is harder for firms to implement them.

More formally, because firms have private information, they play signaling games with uninformed consumers. These signaling games feature pooling equilibria – where the monopolist posts the same price independently of the realization of the shock – and also separating equilibria – where the monopolist posts a different price depending on the realization of the shock. If the proportion of informed is low, separating equilibria do not mimick perfect information allocations and prices – they are distortionary. If the proportion of informed consumers is low enough, pooling equilibria exist as well. As a way of comparison, which type of equilibria delivers higher (ex-ante) profits to the monopolist depends on the proportion of informed consumers. If the proportion of informed is low, pooling equilibria deliver higher profits. This result holds even if firms' nominal costs increase when there is a positive monetary shock.

Second, I show that whether firms adjust prices or not, by having an impact on consumer learning, they inflict an informational externality on other firms. In fact, when consumers see a price change, they change their beliefs about the state of the world, and this has an impact on the other firms' cost of changing prices. For instance, firms playing pooling equilibria prevent consumers from learning, and therefore make it hard to other firms to change prices. This lack of consumer learning hurts firms' profits, and at the same time reduces consumer welfare. Interestingly, the degree of price rigidity depends on firms' marginal costs, and there exist policies capable of improving welfare by bringing the economy closer to perfect information allocations.

More precisely, I show that a temporary tax on monopolists reduces the amount of rigidity, leading to faster learning among consumers and higher welfare.

Third, the dynamic responses of output and inflation are delayed, hump-shaped, and asymmetric. In the model, whether firms adjust to prices or not depends on firms' characteristics, such as firm specific productivity. Therefore, if firms are heterogeneous, some firms adjust prices before other firms. Price changes allow consumers to learn, and consumer learning makes other firms adjust prices in future periods. This reinforcement effect between learning and price changes leads, for a large set of parameter values, to a hump in the response of inflation, because an increasing number of firms change prices at later stages of the adjustment process.

The model also delivers hump-shaped responses of output. The procyclical effect of money arises each time an informed consumer buys from a firm playing a pooling equilibrium, for the usual reason: the informed knows that the price has not adjusted and therefore he buys, in the case of a positive (negative) monetary shock, more (less) than in the benchmark case. This means that two forces drive the size of the aggregate effect: the proportion of informed consumers, and the proportion of firms playing a pooling equilibrium. Initially, as long as the initial proportion of informed consumers is small, the effect on output is small. As the fraction of informed consumers increases, this effect grows. However, over time, as the number of firms playing a pooling equilibrium decreases, the procyclical effect of money dies out, leading to money neutrality in the long run[2]. The asymmetry of the responses of output and inflation relies on the fact that, in the aggregate, prices adjust faster to positive than to negative shocks, consistent with the evidence in Peltzman (2000). As a consequence, negative shocks trigger larger output responses than positive shocks. To summarize, the model delivers delayed, hump-shaped and asymmetric responses of both output and inflation. The first two characteristics seem consistent with the data according to Christiano, Eichenbaum, and Evans (2005), and the third according to Cover (1992).

Related literature. This paper is related to several strands of the literature. First, my paper is complementary to other explanations of the real effects of money through imperfect information (Lucas 1972; Woodford 2002; Mankiw and Reis 2002; Mackowiak and Wiederholt 2009b; Woodford 2009; among others). Most of these papers develop models where firms are imperfectly informed. In this literature, the most closely related papers to mine are Reis (2009) and Mackowiak and Wiederholt (2009a). Both of these papers develop important DSGE models where both firms and households have limited information. The mechanism generating rigidity in these papers is that, because of imperfect information, firms fail to recognize early enough that they should change prices. My model uses a different mechanism for nominal rigidities,

[2] A similar dynamic result is obtained by Mackowiak and Wiederholt (2009a).

based on the idea that at some point firms have better information than the most uninformed portion of consumers[3]. Hellwig and Venkateswaran (2009) develop a model with aggregate and firm-specific shocks. Firms learn from demand, and interestingly this leads to transient effects of money. My setup is more restrictive because I do not allow for firm-specific or idiosyncratic shocks, which facilitates learning on the side of firms. However, this allows to focus on the type of rigidity I want to highlight.

Second, my model is closely related to other economies with decentralized markets and private information (Golosov, Lorenzoni, and Tsyvinski 2009; Duffie and Manso 2007; Amador and Weill 2007). In these models, information about asset values spreads slowly among a population of traders. In a similar spirit, in my economy information about the aggregate state spreads slowly among firms and consumers.

Third, my paper is related to the literature on herding behavior. In that literature economies feature informational externalities in which information gets "trapped", leading to inefficient outcomes (Banerjee 1992; Bikhchandani, Hirshleifer, and Welch 1992; Chamley and Gale 1994; Chari and Kehoe 2003). A paper closely related to mine is Gorodnichenko (2009). This paper presents an interesting and complementary model where informational externalities also play an important role in delaying learning after monetary shocks. My paper differs in the basic mechanism leading to nominal rigidity.

Fourth, my paper is also related to other explanation of price stickiness based on "strong" reactions among consumers to price changes or strategic concerns. Rotemberg has a series of papers based on the idea of consumer anger (Rotemberg 2005), fairness (Rotemberg 2008), or other behavioral considerations (Rotemberg 2007) that create stickiness. Nakamura and Steinsson (2005) develop an interesting repeated game where, due to consumers' habit formation, firms need to commit to keep their price unchanged for long periods of time. Kimball (1995) explores a model in which demand curves are more elastic for price increases than for price decreases. Koszegi and Heidhues (2005) has shown that if consumers perceive price increases as a loss, firms choose a distribution of prices that exhibits stickiness. Maskin and Tirole (1988) have shown that oligopolistic competition also leads to price rigidities. My paper targets the same time of "consumer-based" origin for rigidities, but it differs with previous explanations in that the friction arising in my model is purely informational.

Finally, my paper is related to the industrial organization literature on prices as signals of quality (Wolinsky 1983; Bagwell and Riordan 1991). In those models, buyers are unsure about the quality of products, and therefore the price is a useful signaling device. In my model,

[3]At a conceptual level, my framework can be thought as an opposite model, in terms of agents' information sets, to Lucas (1972). In Lucas' framework, sellers do not know the price level, but buyers do. Here, the opposite happens: sellers know the price level, but some buyers do not.

uninformed buyers are uncertain about the value of goods with respect to goods bought in future periods. Firms, due to their superior information about the state of the economy, have information about future nominal prices and therefore they can signal it. A related interesting paper here is by Kamenica (2008), who presents a model where consumers make inferences on the basis of product lines.

The paper is organized as follows. In Section 2 I present a simple static model that shows the forces at play every time a firm and a group of consumers meet in the dynamic setup. This static model will illustrate how equilibria with nominal rigidity arise, and characterize them. In Section 3 I present the main dynamic results using a simple framework. I show that the fact that firms do not change prices inflicts an externality on other firms. I characterize welfare improving policies capable of correcting this externality. I then simulate the dynamic responses of the economy to a monetary shock. Section 4 develops a cash in advance dynamic model, to show that the insights of the dynamic model can be applied to a monetary framework. Section 5 concludes. Most of the proofs are relegated to the Appendix.

2 A Simple Static Model

This simple static model is meant to provide intuition for the main mechanisms at play in the more involved dynamic framework of the next Section. I consider the problem of a monopolist selling to uninformed consumers. These consumers buy first the good supplied by the monopolist, and then buy all other consumption. The price of all other consumption is exogenous, and can be thought as determined by the long-run macroeconomic state. For instance, after a monetary expansion, these (nominal) prices increase. Uninformed consumers do not know the state, implying that they do not know the price of all other consumption, and thus are uncertain about how much to buy for a given monopolist's price. The main result is that this can lead to rigidity in the monopolist's price.

More specifically, the setup is as follows. A monopolist sells to a unit mass consumers. Consumers are heterogeneously informed about the price of other goods in the economy, which I refer to as the price level. More specifically, consumers are either informed or uninformed. Informed consumers know the price level, uninformed consumers do not know the price level. To simplify, in this stark static model, the monopolist is informed and knows the price level[4]. All consumers observe the price posted by the monopolist before making a purchase decision. Because of this, uninformed consumers are able to update their beliefs about the price level upon observation of the monopolist's price. The main goal of this section is to show that the game between the firm and uninformed consumers has pooling equilibria, where the monopolist's

[4]This assumption will be relaxed in the dynamic model.

price does not depend on the price level, and that these pooling equilibria dominate separating equilibria, i.e. equilibria where the monopolist's price changes with the price level.

Consumers. There is a unit mass of consumers indexed by i. Consumer i has the following utility function of consumption:

$$u(c_i) + C_i \quad . \tag{1}$$

I make the following assumptions concerning the utility function $u(c_i)$.

Assumption 1 *The utility function $u(c_i)$ is assumed to be twice continuously differentiable on R^2_{++}, increasing, and strictly concave.*

The budget constraint is

$$pc_i + PC_i = \text{Income}_i \quad . \tag{2}$$

Utility (1) is linear in C_i and therefore consumption of good c is independent of income. A possible interpretation is that spending on c is a small proportion of total income. Under this interpretation, I refer to good c as a particular consumption good, and to good C as all other consumption of the individual, and to its price P as a "price level".

Goods c and C are bought sequentially. The consumer first buys good c. Then, the consumer buys all other consumption C. Consumers buy good c from a monopolist who sets the price p. P is drawn from a binary probability distribution over $\mathfrak{P} = \{P^h, P^l\}$, where $P^h > P^l$. I refer to $P = P^h$ as the high state, and to $P = P^l$ as the low state. I assume that both states are equally likely: $Pr(P = P^h) = Pr(P = P^l) = 1/2$. Income_i is consumer i's income.

Information. Informed consumers know the realization of the price level when buying from the monopolist. There is a proportion α of informed consumers. The complementary proportion $1 - \alpha$ of consumers is uninformed and does not know the price level when buying from the monopolist. Uninformed consumers know the distribution of possible realizations. All consumers know their income when buying from the monopolist. Also, all consumers observe the price set p when deciding how much to buy from the monopolist.

The monopolist is informed, i.e., he observes the price level before setting his price. Also, to simplify the analysis, it is assumed the monopolist knows the proportion of informed consumers. However, the monopolist cannot discriminate between informed and uninformed consumers.

Informed consumers maximize (1) subject to (2) under perfect information. These consumers know P and maximize their utility without uncertainty.

The monopolist and uninformed consumers play the following one-shot game. First, the monopolist observes the realization of the price level P. After having observed the price level, the monopolist posts a price p. Uninformed consumers observe p, form beliefs μ about the price level, and decide how much to demand from the monopolist[5]. Formally, this sequence of events define a signaling game. The sender of the signaling game is the monopolist. The type of the sender is defined by referring to different possible information sets he can access[6]. Therefore, there are two possible types of monopolist: the monopolist who observes a high realization of the price level, P^h, and the monopolist who observes a low realization of the price level, P^l. The message of the sender is the price p. The receiver is the set of uninformed consumers, whose action is their demand $c_i(\cdot)$, where i belong to the set of uninformed. This action depends on beliefs μ_i.

Monopolist's Problem. The monopolist is able to produce c at zero cost. The monopolist chooses p to maximize revenues:

$$\max_p pc(p, P, \mu_i(p)) \tag{3}$$

where $c(\cdot)$ is total demand for good c, to be derived below. The monopolist sets a price p. Consumers observe the price p and submit their demand. Then, production takes place, and the monopolist sells as much as it is demanded. As it will become clear, total demand $c(p, P, \mu_i(p))$ depends on three objects. First, it depends directly on the price p. Second, it depends on the price level P, because the demand of informed consumers depends on P. Third, it depends beliefs held by the uninformed $\mu_i(p)$, which in turn depend on the monopolist's price p.

Equilibrium definition. I now define a perfect Bayesian equilibrium of the game between the monopolist and the uninformed consumers. I first describe the strategy of the monopolist. I focus on pure strategies. A pure strategy for the monopolist p is a mapping

$$p : \mathfrak{P} \longrightarrow R_+ \quad , \tag{4}$$

that assigns a price p to each state of nature $P \in \mathfrak{P}$. Next, I describe beliefs $\mu_i(p)$ of a given uninformed consumer i. I focus on symmetric beliefs of the uninformed. Beliefs are a probability distribution over $\mathfrak{P}$ defined by a mapping

$$\mu_i : R_+ \longrightarrow [0, 1] \quad , \tag{5}$$

[5]An interesting related paper here is by Jones and Manuelli (2002), where both the informed buyer and informed seller problem is considered.

[6]This is the standard definition of "type" in game theory.

that assigns a probability $\mu_i(p)$ to the high state of nature P^h. The mapping (5) is consistent with Bayes' rule on the path of equilibrium play. Because I focus on pure strategies for the monopolist, the requirement is simply that, for any equilibrium prices (4), denoted $p(P^h)$ and $p(P^l)$ for the high and low states respectively, if $p(P^h) \neq p(P^l)$ (a separating equilibrium), then $\mu_i(p(P^h)) = 1$ and $\mu_i(p(P^l)) = 0$. If instead $p(P^h) = p(P^l)$ (a pooling equilibrium), then $\mu_i(p(P^h)) = \mu_i(p(P^l)) = 1/2$. Beliefs $\mu_i(p)$ are unrestricted for other prices.

I now describe the strategy of uninformed consumers. I focus on symmetric pure strategies. A symmetric pure strategy c_i for a given uninformed consumer i is a mapping

$$c_i : R_+ \times [0,1] \longrightarrow R_{++} \quad ,$$

that assigns a demand c_i to each price p and beliefs $\mu_i(p)$. A perfect Bayesian equilibrium requires that both the firm and the uninformed consumers play a best response. Given this definitions, I can now define an equilibrium formally.

Definition 1 *A Perfect Bayesian Equilibrium (PBE) is a list ($p(P)$,$\mu_i(p)$,c_i), for all i, such that*

1. *There is no profitable deviation from posting p, given consumers' play,*
2. *$\mu_i(p)$ is derived using Bayes' rule on the equilibrium path,*
3. *consumption decisions c_i maximize utility (1), given the budget constraint (2), beliefs $\mu_i(p)$ and firm's play.*

Having defined an equilibrium of the game, it is now useful to present consumers' optimality conditions for good c. This will provide intuition for the fundamental strategic tension between the monopolist and uninformed consumers.

Consumers' Optimality Conditions for good c. In the case of informed consumers, marginal utility of c_i is equated to the relative price of the goods:

$$u'(c_i) = \frac{p}{P} \quad . \tag{6}$$

This equation (6) pins down the demand for good c by consumer i:

$$c_i\left(p\frac{1}{P}\right) \quad . \tag{7}$$

Notice that because of the quasilinearity of preferences, the demand (7) does not depend on income.

In the case of uninformed consumers, marginal utility of c_i is equated to the expected relative price of the goods[7]:

$$u'(c_i) = E_{\mu_i(p)}\left[\frac{p}{P}\right] \quad , \tag{8}$$

where $E_{\mu_i(p)}\left[\frac{p}{P}\right]$ is simply the expectation of relative prices using beliefs $\mu_i(p)$, that is

$$E_{\mu_i(p)}\left[\frac{p}{P}\right] = \mu_i(p)\frac{p}{Ph} + (1-\mu_i(p))\frac{p}{Pl} \quad .$$

Because the expectation $E_{\mu_i(p)}\left[\frac{p}{P}\right]$ is conditional on p (consumer i observes it), the price p can be taken out of the expectation operator, to obtain the demand function

$$c_i\left(pE_{\mu_i(p)}\left[\frac{1}{P}\right]\right) \quad . \tag{9}$$

At this point, it is important to notice that both the demand of the informed (7) and the demand of the uninformed (9) depends on the price chosen by the monopolist p times a deflator. In the case of the informed, this deflator is equal to the inverse of the price level $1/P$. In the case of the uninformed, this deflator is equal to a belief about the inverse of the price level $E_{\mu_i(p)}[1/P]$. Since I focus on symmetric strategies for uninformed consumers, I write total demand as

$$c(p,P,\mu_i(p)) = \alpha c_i\left(p\frac{1}{P}\right) + (1-\alpha)c_i\left(pE_{\mu_i(p)}\left[\frac{1}{P}\right]\right) \quad . \tag{10}$$

An interesting feature of (10) is that is it increasing in $\mu_i(p)$. This fact will be key for the form of equilibria under imperfect information. It implies that the firm prefers uninformed consumers to believe that the price level is high, because in that case the deflator $E_{\mu(p)}\left[\frac{1}{P}\right]$ is low. Notice that this fact is independent of the state of the world.

2.1 Perfect Information Benchmark

To develop intuition, here I consider a perfect information benchmark where all consumers are informed, i.e. when $\alpha = 1$. In this case, total demand is

$$c(p,P,\mu_i(p)) = c_i\left(p\frac{1}{P}\right) \quad . \tag{11}$$

Plugging (11) into (3), the monopolist's problem becomes

[7] To get this expression, substitute Income$_i$ from (2) into (1) and then take the first order condition with respect to c_i.

$$\max_p pc_i\left(p\frac{1}{P}\right) \quad . \tag{12}$$

Lemma 1 *When all consumers know the value of the price level, the monopolist's price is proportional to the price level and demand is the same in both states of nature.*

Proof. Taking the first order condition for the problem (12) and rearranging, get

$$c_i\left(p\frac{1}{P}\right) + p\frac{1}{P}c_i'\left(p\frac{1}{P}\right) = 0 \quad . \tag{13}$$

From condition (13) it is easy to see that the monopolist's optimal price is proportional to the price level. Since total demand depends on the monopolist's price times the realization of the price level, this demand will be the same in both states of nature.

■

2.2 Heterogeneous Information

In this case, there is a proportion of consumers that do not know the realization of the price level P, and therefore $\alpha < 1$. Here I will analyze the equilibria of the game between the firm and uninformed consumers[8]. I will show that, for low enough values of α this game admits equilibria with price rigidity, i.e. pooling equilibria in which the firm posts the same price in both states of the world. I will also show that, for low enough values of α, these equilibria are deliver higher (ex-ante) profits for the monopolist.

This signaling game belongs to the well-known class of monotonic signaling games (Cho and Sobel 1990) (the proof is in Appendix A.1). This has two implications. First, the firm is better off if uninformed customers believe that the state of the world is high, independently of the actual realization of the state. This characteristic follows from the fact that total demand and profits are decreasing in uninformed consumers' beliefs, as explained above. Second, the game has the single-crossing property. This means that the high type is more at ease in posting high prices than the low type. The reason is that informed consumers know that he is the high type,

[8] An interesting, non-strategic, feature of this model without marginal costs is worth highlighting. Suppose no consumer is informed ($\alpha = 0$), and fix consumers' beliefs $E_{\mu_i(p)}\left[\frac{1}{P}\right] \equiv \Xi$. Then, to maximize its revenue, the monopolist posts a price p inversely proportional to consumers' beliefs. To see this, take the first order condition

$$c_i\left(p\Xi\right) + p\Xi c_i'\left(p\Xi\right) = 0 \quad ,$$

which implies that p is inversely proportional to Ξ. The optimality result of pooling equilibria for few informed consumers presented later in this section hinges on this fact, but is valid even in the presence of marginal costs that are proportional to the price level P.

and their demand is less sensitive to price increases compared to the demand from the low type's informed consumers. Together, these two properties make this game tractable.

Signaling games usually have many equilibria, and this game is no exception. I will first characterize separating equilibria. The following lemma characterizes a benchmark separating equilibrium, the one where both types get the highest (ex-post) profits possible.

Proposition 1 (Best Separating Equilibrium) *The following is the Best Separating Equilibrium. Define $\overline{\alpha}$ by*

$$p^l c_i\left(p^l \frac{1}{P^l}\right) = p^h\left(\overline{\alpha} c_i\left(p^h \frac{1}{P^l}\right) + (1-\overline{\alpha})c_i\left(p^h \frac{1}{P^h}\right)\right) \quad . \tag{14}$$

Then,

- *If $\alpha \geq \overline{\alpha}$:*
  - *The firm posts the same prices as in the perfect information benchmark:*

$$p^h = \arg\max_p pc_i\left(p\frac{1}{P^h}\right) \quad , \tag{15}$$

$$p^l = \arg\max_p pc_i\left(p\frac{1}{P^l}\right) \quad . \tag{16}$$

  For a given equilibrium set of prices $p(P)$, define ex-ante real profits as

$$\Pi(p(P)) = \frac{1}{2}\frac{1}{P^h}\pi(P^h) + \frac{1}{2}\frac{1}{P^l}\pi(P^l) \quad . \tag{17}$$

  where $\pi(P) = pc(p, P, \mu_i(p))$. In this case, ex-ante real profits $\Pi(p(P))$ are equal to ex-ante real profits in the perfect information benchmark:

$$\Pi^* = \frac{1}{2}\frac{1}{P^h}\pi(P^h) + \frac{1}{2}\frac{1}{P^l}\pi(P^l) \quad , \tag{18}$$

  where $\pi(P^h) = \max_p pc_i\left(p \cdot 1/P^h\right)$ and $\pi(P^l) = \max_p pc_i\left(p \cdot 1/P^l\right)$.

- *If $\alpha < \overline{\alpha}$:*
  - *The firm posts p^l and $\overline{p} > p_h$ s.t.*

$$p^l c_i\left(p^l \frac{1}{P^l}\right) = \overline{p}\left(\alpha c_i\left(\overline{p}\frac{1}{P^l}\right) + (1-\alpha)c_i\left(\overline{p}\frac{1}{P^h}\right)\right) \quad . \tag{19}$$

  In this case, $\overline{p}$ is strictly decreasing and, if $u(c_i)$ is such that the profit function is single-peaked, then $\Pi(p(P))$ is strictly increasing in α.

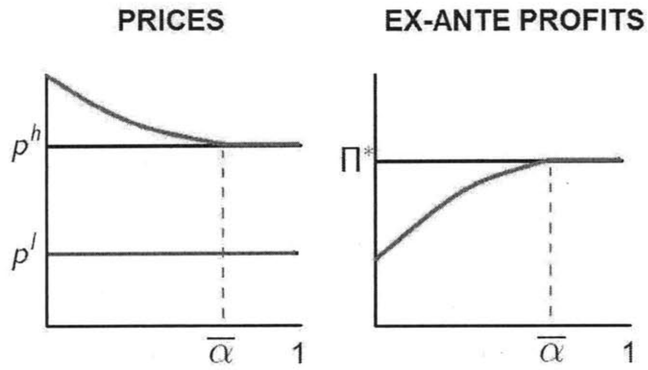

Figure 1: Best Separating Equilibrium (proposition 1) in the case of a quadratic $u(c_i)$.

The proof of this proposition is in the appendix. This lemma shows that when the proportion of informed consumers is high enough, the high type can separate from the low type by posting the perfect information prices. The reason is that, in this case, the proportion of informed consumers is high enough to discourage the low type from imitating him: if the low type posts p^h, the informed know that his price is too high and they buy less than the perfect information quantity c^{ss}, hurting the low type's profits. When the proportion of informed is lower, the only way a separating equilibrium possible is one where the high type posts a price strictly higher than p^h. This makes sure that the low type does not imitate him. Figure 1 is a graphical illustration of this proposition for the case of a quadratic utility function $u(c_i)$. It is shown that ex-ante profits are increasing in α, and reach Π^* when $\alpha \geq \overline{\alpha}$.

Having characterized a benchmark separating equilibrium, I will now characterize a benchmark pooling equilibrium. Pooling equilibria are interesting for the study of nominal rigidities since in these equilibria the firm sets the same price independently of the state of the world. Pooling equilibria exist when the proportion of informed is low. The following proposition characterizes a benchmark pooling equilibrium. This is the pooling equilibrium at the price corresponding to profit maximization when α is equal to zero and no consumer knows the state of the world. In the dynamic model of Section 4, this equilibrium corresponds to keeping the price unchanged after a monetary shock. I also show that when α is equal to zero, this equilibrium reaches the perfect information level of ex-ante profits Π^*[9].

Proposition 2 (p^*-pooling Equilibrium) *Consider p^* s.t.*

$$p^* = \arg\max pc_i \left(p \left[\frac{1}{2} \cdot \frac{1}{P^h} + \frac{1}{2} \cdot \frac{1}{P^l} \right] \right) \quad , \tag{20}$$

[9]This holds also in the presence of marginal costs proportional to the price level P.

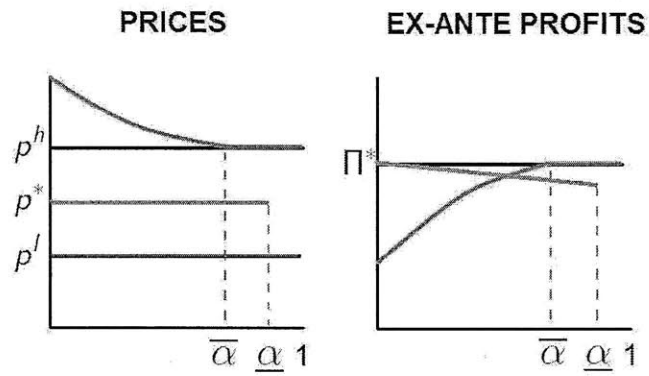

Figure 2: Best Separating Equilibrium and p^*-pooling Equilibrium (Propositions 1 and 2) in the case of a quadratic $u(c_i)$.

and consider $\underline{\alpha}$ s.t.

$$p^*\left(\underline{\alpha}c_i\left(p^*\frac{1}{P^h}\right)+(1-\underline{\alpha})c_i\left(p^*\left[\frac{1}{2}\cdot\frac{1}{P^h}+\frac{1}{2}\cdot\frac{1}{P^l}\right]\right)\right)=\arg\max_p p\left(\underline{\alpha}c_i\left(p\frac{1}{P^h}\right)+(1-\underline{\alpha})c_i\left(p\frac{1}{P^l}\right)\right) \quad (21)$$

For $\alpha \leq \underline{\alpha}$, there exists a pooling equilibrium at p^*. Moreover, if $\alpha = 0$, this equilibrium reaches ex-ante profits Π^*. Also, if $u(c_i)$ is such that the profit function is single-peaked, ex-ante profits $\Pi(p^*)$ are strictly decreasing in α.

The proof is in the appendix. When the proportion of informed consumers is low this equilibrium exists because both types do not want to deviate. In the general equilibrium model this pooling equilibrium will be particularly interesting because, for a strictly positive proportion of informed consumers, monetary shocks will have a procyclical effect on demand. Figure 2 is a graphical illustration of this proposition for the case of a quadratic utility function $u(c_i)$. It is shown that ex-ante profits are increasing in α, and that there is a unique α^{***} where the ex-ante profit functions cross.

An interesting feature of this pooling equilibrium is that it yields higher ex-ante profits for the firm when the proportion of informed consumers is low than any separating equilibrium. The next proposition develops this result in the case of any utility function satisfying assumption 1.

Proposition 3 *If $\alpha \leq \alpha^*$, the p^*-pooling Equilibrium delivers higher ex-ante profits than any separating equilibrium.*

The proof is in the appendix. The intuition for this result is the following. There is an ex-ante trade-off between two possible distortions. The first distortion arises in separating equilibria: in any separating equilibrium, when the proportion is informed consumers is low enough, there is a distortion at the top, because the firm needs to post a very high price to be able to credibly signal the state of the world to uninformed consumers. This distortion hurts ex-ante profits. On the other hand, another type of distortion arises in pooling equilibria: in any pooling equilibrium the price posted does not correspond to beliefs of the informed in all states of nature, making them buy either too much or little with respect to c^{ss}, the demand under perfect information, creating a distortion that hurts ex-ante profits. The first type of distortion is bigger the lower the proportion of informed consumers: in this case, it is easy for the low type to imitate the high type, and therefore the high type needs to post very high prices to separate. The opposite happens in the second type of distortion: this distortion is bigger the higher proportion of informed consumers, in this case the proportion of sales at a wrong price being important. Thus, each of these distortions varies monotonically but in opposite directions with the proportion of informed consumers, creating a trade-off for the firm, which favors pooling equilibria when few consumers are informed, and separating equilibria when many consumers are informed. As shown in the appendix, this holds even in the presence of marginal costs proportional to the price level P.

A symmetric result follows when the proportion of informed consumers is high enough.

Proposition 4 *If $\alpha \geq \alpha^{**}$, the Best Separating Equilibrium delivers higher ex-ante profits than any pooling equilibrium.*

The proof is in the appendix. The intuition for this result is the same as for the previous result. There is a trade-off between two types of distortions. The distortion arising in the Best Separating Equilibrium is small when the proportion of informed consumers α is high, and therefore in this case this equilibrium ex-ante dominates any pooling equilibrium.

In order to apply this framework to a dynamic study of the responses to a monetary shock, an issue that I need to confront is equilibrium selection. Given propositions 3 and 4, it is tempting to pick, for each α, the equilibrium that provides maximum ex-ante real profits to the firm. Thus, when I study these dynamic responses in the next Section (p. 20) I select, for each proportion of informed α, among the Best Separating Equilibrium and the p^*-equilibrium, the one in which firms have higher ex-ante profits. This selection criterion can be justified by assuming that the firm has a commitment device that allows to choose, before the state is realized, a pricing plan. However, these pricing plans have to be 'credible', i.e. the consumers need reasons to believe that, ex-post, the firm will keep its promise. In other words, the firm can only commit to prices that satisfy a PBE and therefore, when the state of the world is low, the firm will not be tempted to fake that the state of the world is high[10]. This type of criterion is useful for my purposes given that for the macroeconomic application of this simple model I am specially interested in

[10] A formal way of justifying this criterion is to modify the game as in Maskin and Tirole (1992).

pooling equilibria, and a particular pooling equilibrium – the p^*-pooling equilibrium – ex-ante dominates all separating equilibria when the proportion of informed is low[11].

However, it is important to emphasize that the above proposed criterion is not crucial for any of the dynamic results of Section 3, including the informational externality and the general shape of the responses to a monetary shock. As long as a non-zero fraction of firms play a pooling equilibrium, similar results would be obtained. What is crucial is that, for low proportion of informed α, the pooling equilibrium exists, and that for high values of α only separating equilibria exist, so that when enough consumers become informed all firms change prices and the economy returns to steady state. Thus, any selection criterion that allows to have a non-zero fraction of firms play pooling equilibria would work[12].

2.3 Comparative Statics in the Presence of Marginal Costs

In a more general model, all the cutoffs presented above should depend on firm specific characteristics. To illustrate this point, let me consider the case where the monopolist has a marginal cost of production kP. I analyze which among the Best Separating vs. the p^*-pooling Equilibrium are ex-ante optimal. The following numerical result follows.

Result 1 (Optimality of p^*-pooling vs. Best Separating Equilibria under Marginal Costs) *Assume $u(c_i) = ac_i - \frac{1}{2}c_i^2$, and consider the Best Separating Equilibrium and the p^*-pooling equilibrium. For $k \leq \hat{k}$, there is $\alpha \leq \hat{\alpha}$ where both equilibria exist. In this region:*

- *for $\alpha \geq \alpha^{***}(k)$, ex-ante profits are higher in the separating equilibrium,*
- *for $\alpha < \alpha^{***}(k)$, ex-ante profits are higher in the pooling equilibrium.*

Moreover, $\alpha^{***}(k)$ is decreasing with k.

As this result shows, which equilibrium delivers higher ex-ante profits depends on firms' marginal cost kP. The higher k, the lower the critical value of α at which firms prefer playing separating equilibria. Figure 3 plots this cutoff as a function of k and shows that it is decreasing. The region below the curve is where pooling equilibria deliver higher ex-ante profits, then region above the curve is where separating equilibria deliver higher ex-ante profits. The intuition is that higher marginal costs need to create less of a distortion in separating equilibria.

[11]In the literature, there is no consensus on how to select equilibria in signaling games. A popular criterion is the intuitive criterion. However, in a richer model with more states of the world – clearly a relevant extension of this model for the analysis of monetary policy – this criterion looses its bite – in the sense that it fails to select a unique equilibrium – and therefore would not be useful (Cho and Kreps 1987, p. 212).

[12]For instance, if one were to Pareto rank equilibria, this would also favor pooling equilibria for low αs and separating equilibria for high αs. The reason is that pooling equilibria are less distortionary in the former case – and thus provide higher welfare, and separating equilibria in the latter.

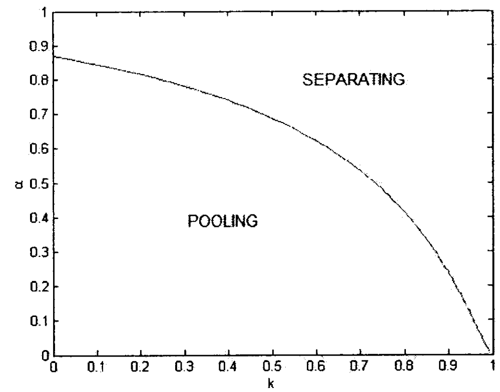

Figure 3: Cutoff $\alpha^{***}(k)$, and regions where p^*-pooling/best separating equilibria deliver highest ex-ante profits ($P^h = 1$,$P^l = .9$,$a = 1$).

This result has an interesting application in a macroeconomic model. Indeed, one can write a model where firms have heterogeneous productivity, giving rise to heterogeneity in k. The presence of firms playing separating equilibria allows for the possibility of consumer learning. This, in turn, has implications for the proportion of firms willing to play separating equilibria, because if the proportion of informed is higher, according to the above result the Best Separating Equilibrium is more profitable when the proportion of informed is higher, or according to Proposition 2, the p^*-equilibrium stops existing. Thus, it seems, a dynamic model can deliver interesting feedback effects between learning and the proportion of firms playing separating equilibria. I explore these themes in the next section.

3 Dynamics Under Heterogeneous Information

In this Section I present and discuss the main dynamic implications of the nominal rigidity developed in Section 2. The discussion will be centered around the informational implications of pooling and separating equilibria. As explained above, uninformed consumers learn the state of the world by observing firms' prices. Therefore, whether firms change prices or not has informational implications for consumers, and therefore, in a dynamic setting, has implications for price setting of other firms. Thus, a given pricing strategy of firms creates an externality on other firms. In particular, firms playing pooling equilibria prevent consumers from learning the state of the world. I first discuss the normative implications and then positive implications of this externality.

3.1 Simple Dynamic Setup

For presentational purposes here I present a simple dynamic setup. This is a simplified version of more involved model of next Section.

Time is discrete, indexed by t, and runs from $t = 1$ to a final period $t = T + 1$. There is a unit mass of consumers. Consumers enter this economy with an initial amount of cash M_0. Consumers buy an initial set of goods $c_1, \ldots, c_T$ using credit, and a good c_{T+1} using cash.

An exogenous monetary shock ν occurs. This shock is a transfer to consumers' cash holdings, whose value is stochastic and observed by consumers only at $T+1$. The sum of the initial amount of cash and the transfer is equal to money supply M:

$$M \equiv M_0 + \nu \quad ,$$

which is high (M^h) or low (M^l), both outcomes equally likely. An exogenous proportion α_1 of consumers learns the amount of the transfer at $t = 1$.

This economy is divided into a unit mass of islands. On every island there is a firm. This firm is a monopolist on the island and sets a price. In this simple dynamic model firms have perfect information and know money supply M. This assumption is relaxed in the more explicit model of Section 4, where an additional initial period is introduced where firms can learn M by observing the amount of sales. This reveals M because of the proportion of informed consumers α_1.

Consumption happens in the following way. Each consumer buys goods $c_1, \ldots, c_T$ sequentially, one after the other, and on a different island each time. More specifically, consumers are divided into representative samples of unit mass. Every period t, each of these samples is sent to a different island for consumption[13].

All firms produce the same good c_t, which is the consumption good of period t. Firm located in island j produces this good with a linear technology of labor

$$c_t = A_j L_t \quad .$$

In the final period $T+1$, consumers receive the cash transfer, and then buy a good C_{T+1} sold by an economy-wide competitive firm. Utility is linear in this good. The price of this good, P, is proportional to total money supply M. Firms hire labor from an economy-wide competitive market at a wage[14] W, also proportional to money supply M.

[13]There is a unit mass of these samples.

[14]In Section 4 I show that, after a normalization, $P = W = M$ in a cash in advance general equilibrium model.

3.2 Informational Externality

I now analyze the normative implications of the nominal rigidity presented in Section 2. Recall that uninformed consumers learn the state of the world by observing firms' prices. Therefore, firms that play a pooling equilibrium will prevent consumers from learning. In this dynamic setting, this inflict an informational externality on firms in the next period. Here I show that this type of informational externality hurts firms' profits and reduces consumers' welfare over periods $t = 1, \ldots, T+1$.

For the purpose of the analysis of the informational externality, I will suppose, without loss of generality, that all firms are identical, thus $A_j = A$ for all j (there is need to relax this assumption to analyze the positive implications of the model, see below). This homogeneity assumption implies that all firms have the same marginal costs of production $kW \equiv \frac{1}{A}W$. Suppose that the initial exogenous proportion of informed consumers α_1 is low enough so that the p^*-equilibrium exists, and that it is played at least on some islands at $t = 1$. The following proposition states that in this case firms' (ex-ante) profits and consumers' (ex-ante) welfare are lower than under perfect information.

Proposition 5 (Welfare Effect of Informational Externality) *Assume $0 < \alpha_1 \leq \underline{\alpha}(k)$ and that at least on one island the p^*-equilibrium is played at time $t = 1$. Then, if the profit function is single-peaked, the sum of firms' ex-ante profits over $t = 1, \ldots, T+1$ is lower than under perfect information. Moreover, if $u(c_i)$ is quadratic, the sum of consumers' ex-ante welfare over $t = 1, \ldots, T+1$ is lower than under perfect information.*

The proof is in the appendix. Intuitively, the profit loss comes from the presence of a non-zero proportion of informed consumers. Firms' prices are not optimal for this segment of the market, leading to a loss in profits. The welfare loss comes from the informed consumers as well: their consumption varies across states, and given the concavity of the utility function $u(x_i)$ and linearity of marginal utility, by Jensen's inequality ex-ante welfare is lower. Perhaps surprisingly, uninformed consumers have the same ex-ante welfare as in the perfect information benchmark. The reason is that the price p^* corresponds to their belief and therefore their consumption bundle is not distorted.

Notice also that in the situation described by Proposition 5, some firms do not change prices and uninformed consumers buying from them do not learn. Since all firms are identical, if no firm changes its price, this is an extreme example in which information gets "trapped" forever and the economy does not return to its benchmark level production even if T is arbitrarily large. This result has the flavor of an information cascade à la Banerjee (1992) or Bikhchandani et al. (1992), where informational "cascades" prevent agents from fully aggregating information in the long run.

Interestingly, in this model nominal rigidity is endogenous and, as shown previously, depends on firms' productivity A and by implication on its marginal cost kW. Thus, it is possible to make prices more flexible using a tax proportional to sales. The tax can increase profits and welfare by allowing uninformed consumers to learn. The following proposition makes this statement precise.

Proposition 6 (Welfare Improving Policy) *Assume $0 < \alpha_1 \leq \underline{\alpha}(k)$ and that at least on one island the p^*-equilibrium is played at period $t = 1$. Then, if T is big enough and the profit function is single-peaked, a social planner can increase the sum of firms' ex-ante profits over $t = 1, \ldots, T + 1$ by implementing a proportional tax rate ψP on firms' sales at period $t = 1$. If $u(c_i)$ is quadratic, the social planner can increase the sum of welfare over $t = 1, \ldots, T + 1$ as well.*

Proof (sketch). Consider ψ financed lump-sum s.t. pooling equilibria do not exist anymore and therefore all firms play a separating equilibrium at $t = 1$. Then, from $t = 2$ on, the economy is in perfect information and ex-ante welfare is maximum. Thus, if T is big enough, this policy improves the sum of welfare over all periods.

■

In the presence of marginal costs of production, the cutoff $\underline{\alpha}(k)$ for existence of pooling equilibria depends on the parameter k (equivalently on productivity A). This cutoff is decreasing in k. A tax ψP increases firms marginal costs and therefore, for high enough ψ only separating equilibria exist. Thus, all consumers learn at $t = 1$. From $t = 2$ on the economy is in perfect information, profits and consumer welfare are higher.

3.3 Learning Dynamics and Responses of the Economy to Monetary Shocks

I now analyze the positive implications of the nominal rigidity presented in Section 2. In particular, I will describe the dynamics of learning, output, and prices. These dynamics are meant to represent the responses of the economy to a monetary shock in an infinite horizon model, as the one developed in the next Section.

One of the issues highlighted when describing the normative implications of the informational externality is that, if the exogenous initial proportion of informed consumers is low, firms are homogeneous and equilibrium selection is such that no firm changes its price, then consumers never learn. This implies that the economy is stuck in an equilibrium in which there is no adjustment back to the perfect information level of production. To solve this issue, and allow for learning for arbitrarily low initial proportion of informed, I introduce firm heterogeneity.

Recall the equilibrium selection criterion proposed in Section 2 – among the Best Separating Equilibrium and the p^*-pooling Equilibrium, pick the one that provides higher ex-ante profits to the firm. Indeed, as shown by result 1, if the utility function is quadratic, according to this criterion the cutoff at which firms change prices depends on firms' productivity. In an economy where firm heterogeneity gives rise to different A_j across firms, low productivity firms change prices for lower proportion of informed consumers. This allows consumers buying from those firms to learn the state, and therefore the proportion of informed consumers is always increasing over time. It is important to notice that the use of the equilibrium selection criterion proposed in the last Section is not crucial for this result, and all other results in this Section. Indeed, it is possible to show that all cutoffs $\overline{\alpha}$, $\underline{\alpha}$, α^*, α^{**}, and α^{***} presented in Section 2 depend on firm productivity, and more generally on other firm specific characteristics. In particular, this is the case for the cutoff of existence of pooling equilibria $\underline{\alpha}$, and thus this equilibrium would exist or not for a particular firm, and this can be used to generate consumer learning. What is important is only that, when it exists, it is not completely ruled out for some firms.

To illustrate consumer learning under heterogeneity, Figure 4 plots the evolution of the proportion of informed consumers (α_t) for a small initial (exogenous) proportion of informed consumers (6 percent). The ratio of money supply M in the two states of the world is equal to $M^l/M^h = .95$, and the number of periods is 13 ($T = 12$). The dynamics of learning are the same in the case of a positive or negative monetary shock. The Figure shows that learning is slow at the beginning, and accelerates, leading, in the case of these parameter values, to a convex curve. Let me now turn to the evolution of the proportion of firms with rigid prices. In a benchmark case where the proportion of firms playing the separating equilibrium is constant, this curve would be concave, as consumers would learn at a constant rate, with fast learning at the beginning, and slow learning towards the end. However, generally, the proportion of firms playing the separating equilibrium is increasing over time. To illustrate, Figure 5 plots the proportion of firms playing the pooling equilibrium. Thus, learning is generally slow at the beginning and it accelerates towards the end, leading to a convex curve. The reason for this is the informational externality in the model. As learning increases, the separating equilibrium becomes more attractive and more firms play it. This increases the rate at which consumers learn. Therefore, there is a positive feedback effect between learning and the proportion of firms playing the separating equilibrium which tends to be convex. This feature is relatively robust to using other parameter values.

The dynamic responses of output the economy are determined by the two forces represented in Figures 4 and 5. Each time an informed consumer buys from a firm playing a pooling equilibrium he creates a procyclical effect of money, i.e. higher output than under perfect information when money supply is high, and lower otherwise. Thus, one the hand, other things equal the higher the

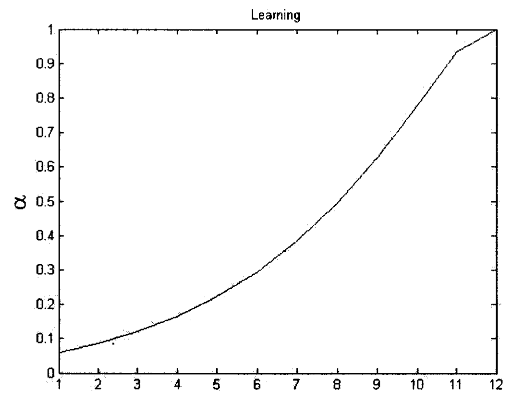

Figure 4: Learning

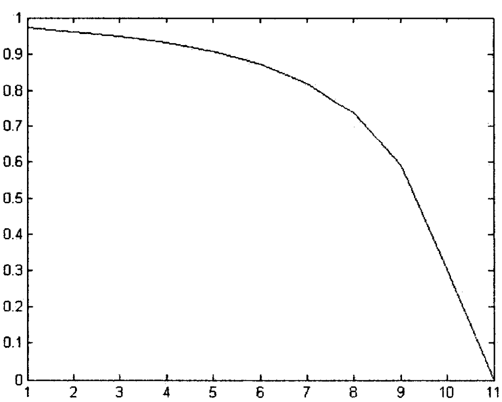

Figure 5: Proportion of firms playing the p^*-pooling equilibrium

proportion of informed, the stronger this effect. On the one other hand, the lower the proportion of firms pooling, the weaker this effect. It is clear that in the long run the fraction firms pooling is zero, and therefore the effect dies out, and the economy returns to the benchmark level of output (informed consumers buying from firms setting perfect information prices do not create a procyclical effect of money). Therefore, as long as there is an initial effect, it will be dampened in the long run. If the initial proportion of informed if small, the effect will be initially and arbitrarily small as well. Thus, learning can create a hump-shaped effect of output, a feature that is shown in Figure 6 as percentage deviations from the perfect information benchmark. This feature is pretty robust to other parameter values. The model is clearly too stylized to be brought to the data, but it is interesting to note the resemblance of these dynamics with its empirical counterparts (Christiano, Eichenbaum, and Evans 2005).

Another feature worth highlighting, regarding the responses of output, is that they are asymmetric – the response is bigger (in absolute value) and arrives a bit earlier when the when the shock is negative. In the case of a negative shock, the effect peaks at around -3 percent, whereas in the case of a positive shock, it peaks at around 2.5 percent.The reason is that sometimes firms increase prices and create a distortion that decreases sales below the benchmark level. This happens when α_t has not yet reached $\overline{\alpha}$, the cutoff for having separating equilibria replicating perfect information (see proposition 1)[15]. This asymmetry in the responses to money shocks is interesting and has been found in the data (Cover 1992).

I now describe the dynamics of prices and inflation. Prices are increasing in the case of high money supply and decreasing in the case of low money supply. This is shown in Figure 7. Notice also that the effect is delayed significantly – as it is in the case of output – because learning is slow at the beginning. This leads to responses of inflation that hump-shaped because, as long as the initial proportion of informed consumers is low enough, the proportion of firms changing prices is increasing over time[16]. More precisely, if the distribution of 'real marginal costs' $k_j = 1/A_j$ is uniform, if the learning curve is convex, the proportion of firms playing a pooling equilibrium has a decreasing slope over time, and there is a hump in inflation. This is the case in this simulation, and this feature is robust to other parameter values. Notice also that the dynamics of prices and inflation are asymmetric. In particular, prices adjust faster in the high state, consistent a response of output that is smaller than in the low state.

To summarize, the model delivers interesting responses of both output and inflation. In the

[15]The reason this effect is small and only leads to a slightly dampened positive shock is that only a few firms cause it, precisely because, according to the selection criterion used firms avoid this distortion. I have experimented with many parameter values and I have never found that this effect dominates, leading positive shocks to have negative effects on output.

[16]This is related to an interesting paper by Cavallo (2009) who, using daily frequency data on price changes, finds empirical evidence that firms synchronize price changes.

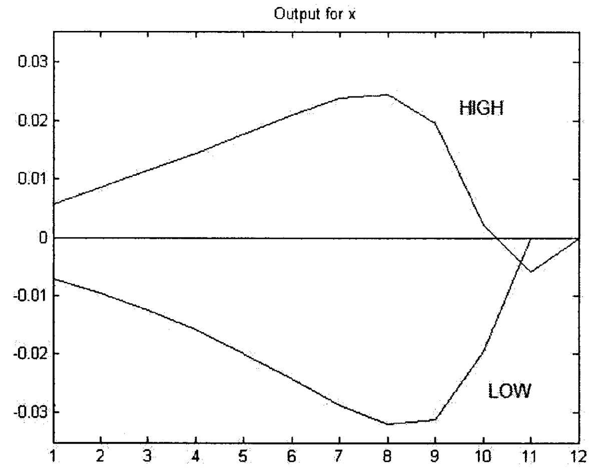

Figure 6: Dynamics of Output in High/Low State

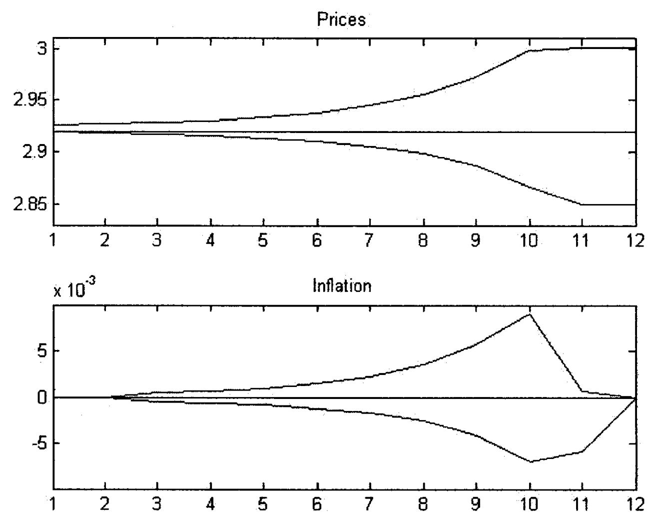

Figure 7: Dynamics of Prices and Inflation in High/Low State

data, these responses have been estimated as delayed and hump-shaped as well. In the model, the informational externality and the implied the slow nature of learning deliver naturally these features, as long as the initial proportion of informed consumers is small.

4 The Dynamic Cash in Advance Framework

The goal of this Section is to show that the simple dynamic model of last Section is compatible with a cash in advance, general equilibrium, framework.

In this model money has an explicit role due to a standard cash in advance constraint (Lucas and Stokey 1987). In each island there is a price-setting firm. Firms are heterogeneous in that they will have different productivity levels. As shown by Result 1, heterogeneity implies that after a monetary shock firms with low productivity change prices earlier. Consumers learn endogenously from these price changes. As consumers become informed, more and more firms change prices, slowly bringing the economy to a long run equilibrium.

Time. Time is discreet and every period indexed by $t = 0, \ldots$. Every period is subdivided into $N + 2$ subperiods, indexed by $\tau = 0, 1, \ldots, N + 1$. For the analysis of the dynamics after a monetary transfer, I am interested in the dynamics across subperiods.

Geography. There is a unit mass of islands indexed by either $j \in [0, 1]$. Every island there is populated by a unit mass of households, indexed by i. Therefore, a household in this economy is identified by a double index ij. On every island, households own a firm that produces a consumption good c. The firm is a monopolist on the island where it operates. Each one of these firms produce the same good.

Shoppers and Workers. Every household is divided into a shopper and a worker.

Consumption. In the first $N+1$ subperiods of every period t shoppers are randomly sent to islands for consumption of goods c_τ. Therefore, in every subperiod $\tau \leq N+1$, every monopolist receives a new mass of shoppers that demand good c_τ. Goods c_τ are sold by the monopolists on credit. For this reason, in the rest of the paper I will refer to these goods as credit goods. In the last subperiod $N + 2$ of every period t, shoppers buy a cash good C from a centralized, representative competitive firm. As in the partial equilibrium model, goods bought in decentralized markets – credit goods – are denoted in lowercase, and goods bought in centralized markets – the cash good – are denoted in uppercase. As it will become clear, goods bought in centralized markets will have prices that will be proportional to the money supply at all times and will be denoted in uppercase. (Goods bought in decentralized markets need not have prices proportional to the money supply at all times, and will be denoted in lowercase.)

Goods and Labor Markets. Goods c_τ are sold on every island by a price-setting monopolist. Good C is sold in a centralized competitive market by a representative firm. There is a centralized competitive labor market where all workers supply labor L. Because labor is supplied in a centralized market it is denoted in uppercase. All firms in the economy hire labor from this market.

Timing and Information Assumptions. At the beginning of every period t there is a monetary transfer ν_t into the economy. I will refer to ν_t as the "monetary transfer" or the "monetary shock" interchangeably. This monetary transfer is not immediately revealed to households.

At the first subperiod $\tau = 0$, every island receives a representative sample of shoppers in the economy. An exogenous proportion of shoppers α_{0t} knows the monetary transfer ν_t when buying credit good c_0. The complementary proportion $1 - \alpha_{0t}$ does not know the transfer when buying the credit good. In this period, firms are uninformed, but can learn the state by observing total demand.

At every subperiod $\tau = 1, \ldots, N$, again, every island receives a representative sample of shoppers in the economy. In the first subperiod $\tau = 1$, the proportion of α_{1t} knows the monetary transfer ν_t when buying credit good c_1. The complementary proportion $1 - \alpha_{1t}$ does not know the transfer when buying the credit good. At this point firms are informed about the monetary shock, and therefore, on some islands it is possible to learn the monetary transfer from prices set by monopolists. Therefore, the proportion of shoppers that know the monetary shock grows endogenously: $\alpha_{2t} \geq \alpha_{1t}$, and more generally, $\alpha_{\tau+1t} \geq \alpha_{\tau t}$, $\tau = 1, \ldots, N$.

Meanwhile, workers supply labor in the centralized labor market. In equilibrium the wage W_t is perfectly flexible and proportional to the money supply, and for this reason, the wage immediately reveals the transfer to workers.

At the last subperiod $\tau = N + 1$, each shopper receives the monetary transfer ν_t and buys the cash good C using cash in a centralized market. At the end of the period, the worker comes back home bringing labor income, the household pays consumption of credit goods, profits from firms are received, and financial markets open. At this point, all agents in the economy know the monetary transfer ν_t. This implies that all agents know the money supply in the economy.

Games Between Monopolists and Uninformed Shoppers. Every time a particular informed firm and a sample of the shoppers population are matched, the firm and the proportion of uninformed shoppers play the following signalling game. Firms' type is determines by the monetary shock ν_t. Therefore, there two possible types for the firm: the "high" type, i.e. the firm that observed a high monetary shock; and the "low" type, i.e. the firm that observed a low monetary shock. After observing the monetary shock, the firm posts a price $p_{j\tau t}$ in a store for the good it sells. Uninformed shoppers observe this price, form beliefs about the monetary shock $\mu(p_{j\tau t})$, and decide how much to demand[17]. Notice that shoppers and firms are matched at every subperiod τt, and also that an infinite number of these matches happen at every τt.

Stochastic Process for Money. Money supply evolves as:

$$\log M_t = \log M_{t-1} + \nu_t \quad , \tag{22}$$

where ν_t is shock that can take two values, $\nu_h > 0$ or $\nu_l < 0$, both outcomes with equal probability. I further impose that

$$E\left[e^{-\nu_t}\right] = 1 \quad . \tag{23}$$

This centering assumption implies that the the inverse of the money supply, which I interpret as the real value of a 1 dollar bill, is a martingale:

$$E\left[\frac{1}{M_t}\right] = E\left[\frac{e^{-\nu_t}}{M_{t-1}}\right] = \frac{1}{M_{t-1}} \quad . \tag{24}$$

This assumption will be convenient when characterizing the equilibria of the games played between the monopolists and the shoppers.

Households' Problem. Household i of island j faces the problem

$$\max E_{ij\tau t}\left[\sum_{t=0}^{\infty}\beta^t\left(\sum_{\tau=0}^{N}u(c_{ij\tau t}) + v(C_{ijt}) - L_{ijt}\right)\right] \quad , \tag{25}$$

where $c_{ij\tau t}$ is consumption of good x at subperiod τ time t, produced by a randomly matched firm $\hat{\jmath}$ of island $\hat{\jmath}$, C_{ijt} is consumption of the credit good, and L_{ijt} is labor supplied by the worker. $\hat{\jmath}(i,j,\tau,t)$ is a function designates the firm $\hat{\jmath}$, that sells to household ij at subperiod τ time t. $E_{ij\tau t}$ is an expectation that is taken at different stages of each period, and using the relevant decision maker's (shopper or worker) specific information set. This maximization is subject to the budget constraint

$$\sum_{\tau=1}^{N} p_{\hat{\jmath}\tau t}c_{ij\tau t} + P_t C_{ijt} + M_{ijt} + B_{ijt} = (1+R_t)B_{ijt-1} + M_{ijt-1} + \tilde{\nu}_t + W_t L_{ijt} + s_{ij}\Pi_{jt} \quad , \tag{26}$$

[17]Notice that all uninformed consumers play the same strategy, and therefore have the same beliefs about the monetary shock.

where $\tilde{\tau}_t$ is a lump-sum transfer consistent with the process of money[18].

The cash-in-advance constraint for consumption of good y is

$$P_t C_{ijt} \leq M_{ijt-1} + \tilde{\nu}_t \quad . \tag{27}$$

At this point it is important to notice that households' preferences are quasilinear. By eliminating wealth effects on household choices, this assumption simplifies the analysis in three ways. First, similar to Lagos and Wright (1995), it allows to handle the heterogeneity in households' information. These heterogeneity causes different shoppers to make different choices for credit goods c_τ. However, the absence of wealth effects implies that all other choices are homogenous across households, and therefore the cross-sectional distribution of money and bond holdings collapses into a degenerate distribution. For this reason, I do not need to keep track of it. A second implication of the absence of wealth effects is that there is no option value in waiting to gather more information before consuming credit goods, which simplifies the games between firms and shoppers enormously. Third, linearity in labor supply implies that every set of equilibria of these games is consistent with a general equilibrium in which the labor market clears.

Monopolists. On every island j there is a monopolist that produces good c_τ using technology

$$c_{j\tau t} = A_j L_{j\tau t} \quad , \tag{28}$$

where $L_{j\tau t}$ is labor hired by firm j at subperiod τt. Firms are heterogenous with respect to their productivity. Write the firm's real marginal costs as

$$k_j = \frac{\frac{W_t}{M_t}}{A_j} \quad , \tag{29}$$

where W_t is the nominal wage. I assume that firm heterogeneity is such that there is a distribution A_j that yields a uniform distribution of real marginal costs

$$k_j \sim U[0, a] \quad . \tag{30}$$

Representative Firm Producing C. There is a representative firm that hires labor and produces good C according to a simple linear technology. The labor productivity if this firm is normalized to 1:

$$C = L \quad .$$

[18]More specifically, $\tilde{\nu}_t$ is such that $\tilde{\nu}_t = M_t - M_{t-1}$. I will show that since in equilibrium all agents have the same money holdings and therefore this way of defining the transfer is possible.

Definition of Equilibrium. A general equilibrium of this economy are allocations $\{c_{ij\tau t}, C_{ijt}\}$, beliefs $\mu_{ij\tau t}(p_{j\tau t})$, labor supply decisions $\{L_{ijt}\}$, prices $\{p_{j\tau t}, P_t\}$, nominal wages, $\{W_t\}$, interest rates $\{1 + R_t\}$, for all i, j, τ, t, s.t.

1. The households' conditions for optimality and associated constraints are satisfied;
2. the equilibrium strategies for the games played between monopolists and shoppers satisfy Bayesian Perfection:
   - monopolists post prices to maximize profits, given consumers' play,
   - uninformed shoppers use Bayes' rule on the path of equilibrium play,
   - shoppers' demand satisfies the condition for optimality;
3. the representative firm maximizes profits taking the price as given;
4. goods, labor, bonds, and money markets clear.

4.1 General Equilibrium.

Household ij's Optimality Conditions. The conditions for optimality are computed as follows. Each time the shopper of household ij is matched with a monopolist, he observed the price that was posted and computes the first order condition for consumption of the good sold by the monopolist:

$$\beta^t u'(c_{ij\tau t}) = p_{\hat{j}\tau t} E_{\mu_{ij\tau t}}[\lambda_{ijt}] \quad . \tag{31}$$

$E_{ij\tau t}[\cdot]$ is an expectation computed using the shopper's information set. This information set contains information he has previously collected plus the information revealed by the price of the monopolist.

When the shopper finally arrives to buy the cash good, he computes a first order condition for consumption of this cash good after observing its price. This good is sold by a representative competitive firm in a centralized market, and therefore its price reveals the realization of the monetary shock to the shopper, in case he did not know it already. Therefore, at this point the shopper does not face any uncertainty, and the first order condition is:

$$\beta^t v'(C_{ijt}) = P_t(\lambda_{ijt} + \psi_{ijt}) \quad . \tag{32}$$

The worker arrives at the market for labor and computes a first order condition for labor supply after observing the equilibrium wage. This is a centralized and competitive market, and therefore this wage reveals the realization of the monetary shock to the worker. Therefore, the worker does not face any uncertainty, and the first order condition is:

$$\beta^t = W_t \lambda_{ijt} \quad . \tag{33}$$

The first order condition for money holdings is computed at financial markets at the end of the period, and therefore under perfect information:

$$\lambda_{ijt} = E_t[\lambda_{ijt+1} + \psi_{ijt+1}] \quad . \tag{34}$$

The first order condition for bond holdings is also computed under perfect information:

$$\lambda_{ijt} = (1 + R_{t+1}) E_t[\lambda_{ijt+1}] \quad . \tag{35}$$

Cash Good and Labor Markets. It is possible to show that every set of PBE of the games played between monopolists and shoppers is consistent with a general equilibrium. The reason is the linearity in the disutility of labor, which implies that individual labor choices are always consistent with the resource constraint of the economy.

In equilibrium, the price of the cash good C is pinned down by the cash in advance constraint, and therefore it is proportional to the money supply. Optimality of production for the representative firm producing the cash good immediately implies that the wage W_t is also proportional to the money supply M_t. After a normalization, we have that all of these three quantities are always equal:

$$P_t = W_t = M_t \quad . \tag{36}$$

Also, the nominal interest rate is determined in expectations of next period's shock, and therefore always constant. Consumption of the credit good is also constant every period.

Demand for Credit Good $c_{\tau t}$ by Shopper ij. Substituting (33) into (31):

$$u'(c_{ij\tau t}) = p_{\hat{j}\tau t} E_{\mu_{ij\tau t}} \left[\frac{1}{M_t} \right] \quad . \tag{37}$$

From this equation get the demand:

$$c_{ij\tau t} \left(p_{\hat{j}\tau t} E_{\mu_{ij\tau t}} \left[\frac{1}{M_t} \right] \right) \quad . \tag{38}$$

As in the partial equilibrium model of section 2, the demand for credit goods depends on a deflated version of the price posted by firms. Here, the deflator is a belief about the inverse of the real value of a unit of money, and therefore the consumer basically transforms the nominal price posted by the monopolist into a real price. If this estimate is low (corresponding to a belief that the supply of money is high), a given nominal price is deflates into a low real price, increasing the shopper's demand. As it will become clear, this effect will limit firms credibility to signal that there has been a high monetary shock in the economy.

Aggregate Demand for $c_{\tau t}$. At every subperiod τt a proportion $\alpha_{\tau t}$ of shoppers know the monetary aggregate. Therefore, the aggregate demand is

$$\int_{\mathcal{M}_{\hat{j}\tau t}} c_{ij\tau t}(p, M_t, \mu_{ij\tau t}) di dj = \alpha_{\tau t} c_{ij\tau t}\left(p_{\hat{j}\tau t} \cdot \frac{1}{M_t}\right) + (1-\alpha_{\tau t}) c_{ij\tau t}\left(p_{\hat{j}\tau t} \cdot E_{\mu_{ij\tau t}}(p_{\hat{j}\tau t})\left[\frac{1}{M_t}\right]\right) \quad , \tag{39}$$

where $\mathcal{M}_{\hat{j}\tau t} \in [0,1] \times [0,1]$ is the subset of consumers matched with firm $\hat{j}$ at subperiod τt. A proportion $\alpha_{\tau t}$ of the shoppers arriving to firm $\hat{j}$ will know the monetary aggregate and therefore will have the correct beliefs, and the complementary proportion will form beliefs using Bayes' rule on the equilibrium path.

At this point, notice that the total demand (39) that every firm faces is the same as (10) in Section 2. This shows that result 1 applies in the case of this cash in advance model[19]. Also, all the dynamic results of Section 3 apply here. Thus, it is possible to write down a cash in advance general equilibrium framework compatible with all the results of this paper.

5 Conclusion

This paper proposes an explanation to nominal rigidities that is complementary to other explanations based on imperfect information[20]. It analyzes an environment where consumers' information about aggregate shocks is heterogenous. Some consumers are informed while others are uninformed. In the model, firms can learn the realization of aggregate shocks by observing total demand. This leads to an asymmetry of information between firms and uninformed consumers. I show that, because of this asymmetry of information, equilibria with rigid prices arise and are, for some parameter values, more profitable to firms than equilibria with flexible prices. In these equilibria firms do not change prices even though they have perfect information and face no menu costs.

This result, together with the strategic tension arising between firms and consumers in my model, accommodate some anecdotal evidence related to price rigidity. For instance, when asked about their reluctance to change prices, often firm managers declare to be concerned about

[19]See page 16.
[20]For a more detailed exposition of the relationship of my paper with this literature, see page 4.

consumers' reactions to price changes (Blinder 1991). Also, it is customary for firms to justify the reasons behind price increases or explicitly apologize once an increase has been implemented.

This type of nominal rigidity has interesting normative and positive dynamic implications. On the normative side, the rigidity leads to an informational externality that hurts firms' profits and reduces consumer welfare. This type of externality has been found on other settings (Banerjee 1992; Chamley and Gale 1994; Gorodnichenko 2009). I show that a temporary tax on monopolists reduces the amount of rigidity and improves welfare. On the positive side, the dynamic responses of output and inflation after a shock are delayed and hump-shaped. The reason is that initially few consumers know about the shock, but learning accelerates as firms start changing prices. These features seem consistent with the data (Christiano, Eichenbaum, and Evans 2005). Moreover, the responses are asymmetric: bigger effects on output for negative shocks, and faster increase in prices after a positive shock. These asymmetric features also seem realistic (Cover 1992; Peltzman 2000).

Further extensions of this setup are needed to understand several issues left open. For instance, a model with a richer type space for monetary policy is a clear extension to understand how the degree of rigidity changes for different shocks. Also, the cash in advanced setup can be further enriched. The first extension that comes to mind is staggering of households so that different set of households make use of financial markets at different times. I leave these extensions for future work.

A Appendix

A.1 Characterization of the Game.

First, it is necessary to define the following well-known property for a function of two variables.

Definition 2 (Increasing Differences Property) *A function $f(x, y)$ has strict increasing differences in (x, y) if, for $x' > x$ and $y' > y$,*

$$f(x, y') - f(x, y) < f(x', y') - f(x', y) \quad . \tag{40}$$

The following lemma shows that under some conditions this game is a standard monotonic signaling game (Cho and Sobel 1990).

Lemma 2 (Characterization of the Game) *If $\alpha > 0$ and $pc_i(p \cdot 1/P)$ has increasing differences in (p, P), this is a monotonic signaling game. It satisfies:*

1. *Monotonicity.*

   Let $\mu_i'(p)$ and $\mu_i(p)$ be two possible beliefs by the uninformed. If $\mu_i'(p) > \mu(p)$, then, for all p, $pc(p, P, \mu_i'(p)) > pc(p, P, \mu_i(p))$.

2. *Single-crossing. For any $p' > p$, we have that, for arbitrary demand by the uninformed, $p'c(p', P^l, \mu_i(p)) \geq pc(p, P^l, \mu_i(p)) \implies p'c(p', P^h, \mu_i(p)) > pc(p, P^h, \mu_i(p))$*

Proof. I first prove monotonicity, and then single-crossing.

1. Monotonicity. $u'(c_i)$ being a strictly decreasing function, the demand of the uninformed $c_i(pE_{\mu_i(p)}[1/P])$ is strictly increasing in $\mu_i(p)$. Therefore, for any $\mu_i'(p) > \mu(p)$, $pc'(p, P, \mu_i'(p)) > pc(p, P, \mu_i(p))$.

2. Single-crossing.

   Consider p, p', such that $p < p'$, and assume

   $$p'c(p', P, \mu_i(p)) \geq pc(p, P, \mu_i(p))$$

   This is equivalent to

   $$p'c(p', P, \mu_i(p)) - pc(p, P, \mu_i(p)) \geq 0$$

   Since $c(p, P, \mu_i(p))$ has strict increasing differences in (p, P),

$$p'c(p', P^h, \mu_i(p)) - pc(p, P^h, \mu_i'(p)) > p'c(p', P^l, \mu_i(p)) - pc(p, P^l, \mu_i'(p)) \geq 0$$

and therefore

$$p'c(p', P^h, \mu_i(p)) > pc(p, P^h, \mu_i(p))$$

■

A.2 Proof of Proposition 1

The cutoff $\overline{\alpha}$ is obtained using the Incentive Compatibility (IC) constraint for the low type (14). Off equilibrium path beliefs are $\mu_i(p) = 0$.

Once the cutoff obtained, there are two cases:

- $\alpha \geq \overline{\alpha}$.

  In this case, (14) is satisfied at p^h and p^l defined by (15) and (16). Firms optimization in each state yields p^h and p^l. Therefore, this is the best separating equilibrium. Since in this outcome consumption is c^{ss}, in this outcome ex-ante profits are Π^*.

- $\alpha < \overline{\alpha}$.

  In this case, the IC constraint for the low type is satisfied for a price $\overline{p}$ defined by (19). In all separating equilibria the low type posts p^l, otherwise he has a profitable deviation. Moreover, from (14) $\overline{p}$ is decreasing in α, and therefore $\Pi(p(P))$ is increasing in α.

■

A.3 Proof of Proposition 2

The cutoff $\underline{\alpha}$ is obtained from the IC constraint for the high type (21). Off equilibrium path beliefs are $\mu_i(p) = 0$. Given this and (21), the IC constraint for the low type

$$p^*\left(\underline{\alpha}c_i\left(p^*\frac{1}{P^h}\right) + (1-\underline{\alpha})c_i\left(p^*\left[\frac{1}{2}\cdot\frac{1}{P^h} + \frac{1}{2}\cdot\frac{1}{P^l}\right]\right)\right) = p^l\left(\underline{\alpha}c_i\left(p^l\frac{1}{P^h}\right) + (1-\underline{\alpha})c_i\left(p^l\frac{1}{P^l}\right)\right) \quad ,$$

is satisfied $\forall\alpha \leq \underline{\alpha}$.

I now show that if $\alpha = 0$, $\Pi(p(P)) = \Pi^*$. For $\alpha = 0$,

$$\Pi(p^*) = \frac{1}{2}\frac{1}{P^h}p^*c_i\left(p^*\left[\frac{1}{2}\frac{1}{P^h} + \frac{1}{2}\frac{1}{P^l}\right]\right) + \frac{1}{2}\frac{1}{P^l}p^*c_i\left(p^*\left[\frac{1}{2}\frac{1}{P^h} + \frac{1}{2}\frac{1}{P^l}\right]\right) \quad . \tag{41}$$

From the FOC of (20), (15) and (16), notice that

$$p^* \left[\frac{1}{2} \cdot 1/P^h + \frac{1}{2} \cdot 1/P^l \right] = p^h \cdot 1/P^h = p^l \cdot 1/P^l \quad . \tag{42}$$

Also,

$$c_i \left(p^* \left[\frac{1}{2} \cdot 1/P^h + \frac{1}{2} \cdot 1/P^l \right] \right) = c^{ss} \quad . \tag{43}$$

Together, (42) and (43) imply that the right hand side of (41) is equal to Π^*.

More generally,

$$\Pi(p^*) = \alpha \left(\frac{1}{2} \frac{1}{P^h} p^* c_i \left(p^* \frac{1}{P^h} \right) + \frac{1}{2} \frac{1}{P^l} p^* c_i \left(p^* \frac{1}{P^l} \right) \right) + (1-\alpha)\Pi^* \quad , \tag{44}$$

which is decreasing in α.

■

A.4 Proof of Proposition 3

In the case of the best separating equilibrium, $\Pi(p(P))$ is continuous and strictly increasing in α. In the case of the p^*-pooling equilibrium, from (44) $\Pi(p^*)$ is continuous at $\alpha = 0$, reaches Π^* at $\alpha = 0$, and it is strictly decreasing thereafter. Thus, there is a boundary $[0, \alpha^*]$ away from $\alpha = 0$ where $\Pi(p(P))$ is higher in the p^*-pooling equilibrium.

■

As argued in the body of the text, this proposition can be extended to the case of marginal costs proportional to the price level P. It is required that these are low enough so that the firm makes profits in the high state of the world. As long as $\alpha = 0$ ex-ante real costs are the same as under perfect information, and there it is very easy to obtain the alleged result.

A.5 Proof of Proposition 4

Pick $\alpha^{**} = \overline{\alpha}$. For $\alpha \geq \alpha^{**}$, $\Pi(p(P)) = \Pi^*$ in the separating equilibrium. In the pooling equilibrium, if it exists for $\alpha \geq \overline{\alpha}$, the right hand side (RHS) of (44) is strictly below Π^*, and decreasing.

■

A.6 Proof of Proposition 5

The RHS of (44) is strictly smaller than ex-ante profits in the perfect information benchmark, $\Pi^*(k)$. This proves the statement on ex-ante profits.

I will now prove the result on ex-ante welfare. Given separability of the utility function, consider the case of good c_1 (the argument is identical in the case of the other goods $c_2, \ldots, c_T$). The demand for this good is

$$c_{i1} = a - p_1^* E_{i1}\left[\frac{1}{P}\right]$$

In case of uninformed, consumption of good c_1 in the p^*-pooling Equilibrium is

$$c_{i1}^h = c_{i1}^l = a - p_1^*\left[\frac{1}{2}\frac{1}{P^h} + \frac{1}{2}\frac{1}{P^l}\right]$$

Because $p_1^*\left[\frac{1}{2}\cdot 1/P^h + \frac{1}{2}\cdot 1/P^l\right] = p^h \cdot 1/P^h = p^l \cdot 1/P^l$, $c_{i1}^h = c_{i1}^l = c^{ss}$. Thus, in the p^*-pooling Equilibrium, the uninformed get the same ex-ante utility as under perfect information.

In the case of the informed, consumption of good c_1 in this equilibrium is either

$$c_{i1}^h = a - p_1^*\frac{1}{P^h} > c^{ss} \quad ,$$

or

$$c_{i1}^l = a - p_1^*\frac{1}{P^l} < c^{ss} \quad .$$

Using Jensen's inequality, for an informed consumer:

$$\frac{1}{2}u(c_{i1}^h) + \frac{1}{2}u(c_{i1}^l) < u(c^{ss})$$

Thus, the informed get strictly less ex-ante utility than under perfect information.

Therefore, welfare is strictly smaller than under perfect information.

■

References

Amador, M. and P. O. Weill (2007). Learning from private and public observation of other's actions. *mimeo, Stanford University.*

Anderson, E. and D. Simester. Price stickiness and customer antagonism. *Quarterly Journal of Economics.* forthcoming.

Bagwell, K. and M. H. Riordan (1991). High and declining prices signal product quality. *American Economic Review 81*(1), 224–239.

Banerjee, A. V. (1992). A simple model of herd behavior. *Quarterly Journal of Economics 57*(3), 797–817.

Bikhchandani, S., D. Hirshleifer, and I. Welch (1992, October). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy 100*(5), 992–1026.

Blinder, A. (1991). Why are prices sticky? Preliminary results from an interview study. *American Economic Review 89-96.*

Blinder, A., E. Canetti, D. Lebow, and J. Rudd (1998). *Asking About Prices: A New Approach to Understanding Price Stickiness.* New York. Russell Sage Foundation.

Cavallo, A. (2009). Scraped data and sticky prices: frequency, hazards, and synchronization. *mimeo, Harvard University.*

Chamley, C. and D. Gale (1994). Information revelation and strategic delay in a model of investment. *Econometrica 62*(5), 1065–1085.

Chari, V. V. and P. J. Kehoe (2003). Hot money. *Journal of Political Economy 111*(6), 1262–1292.

Cho, I.-K. and J. Sobel (1990). Strategy stability and uniqueness in signaling games. *Journal of Economic Theory 50*, 381–413.

Christiano, L. J., M. Eichenbaum, and C. L. Evans (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy 113*(1), 1–45.

Cover, J. P. (1992). Asymmetric effects of positive and negative money-supply shocks. *Quarterly Journal of Economics 107*(4), 1261–1282.

Duffie, D. and G. Manso (2007). Information percolation in large markets. *American Economic Review Papers and Proceedings 97*, 203–209.

Fehr, E., G. Kirchsteiger, and A. Riedl (1993). Does fairness prevent market clearing? *Quarterly Journal of Economics 108.*

Fuhrer, J. C. (2000). Habit formation in consumption and its implications for monetary-policy models. *American Economic Review 90*, 367–90.

Golosov, M., G. Lorenzoni, and A. Tsyvinski (2009). Decentralized trading with private information. *mimeo, MIT*.

Gorodnichenko, Y. (2009). Endogenous information, menu costs and inflation persistence. *mimeo, University of California, Berkeley.*

Greenwald, B. and J. E. Stiglitz (1989). Toward a theory of rigidities. *American Economic Review Papers and Proceedings 79*(2), 364–369.

Hall, R. and C. Hitch (1939). Price theory and business behavior. *Oxford Economic Papers 2*, 12–45.

Heidhues, P. and B. Köszegi (2005). The impact of consumer loss aversion on pricing. *mimeo, University of Bonn and University of California, Berkeley.*

Hellwig, C. and V. Venkateswaran (2009). Setting the right prices for the wrong reasons. *mimeo, UCLA.*

Jones, L. and R. E. Manuelli (2002). Volatile policy and private information: The case of monetary shocks. *Journal of Economic Theory 99*, 265–296.

Kahneman, D., J. L. Knetsch, and R. H. Thaler (1986). Fairness as a constraint on profit seeking: Entitlements in the market. *American Economic Review 76*(4), 728–741.

Kamenica, E. (2008). Contextual inference in markets: On the informational content of product lines. *American Economic Review 98*(5), 2127–2149.

Kimball, M. S. (1995). The quantitative analytics of the basic neomonetarist model. *Journal of Money, Credit and Banking Part 2: Liquidity, Monetary Policy, and Financial Intermediation 27*(4), 1241–1277.

Koszegi, B. and P. Heidhues (2005). The impact of consumer loss aversion on pricing. *Working paper, Berkeley.*

Lagos, R. and R. Wright (1995). A unified framework for monetary theory and policy analysis. *Journal of Political Economy 113*(31), 463–484.

Lucas, R. E. (1972). Expectations and the neutrality of money. *Journal of Economic Theory 4*, 103–124.

Lucas, R. E. and N. L. Stokey (1987). Money and interest in a cash-in-advance economy. *Econometrica 55*(3), 491–513.

Mackowiak, B. and M. Wiederholt (2009a). Business cycle dynamics under rational inattention. *mimeo.*

Mackowiak, B. and M. Wiederholt (2009b). Optimal sticky prices under rational inattention. *American Economic Review 99*(3), 769–803.

Mankiw, G. and R. Reis (2002). Sticky information versus sticky prices: a proposal to replace the New Keynesian Phillips curve. *Quarterly Journal of Economics 117*(4), 1295–1328.

Maskin, E. and J. Tirole (1988). A theory of dynamic oligopoly, II: Price competition, kinked demand curves, and Edgeworth cycles. *Econometrica 56*(3), 571–599.

Maskin, E. and J. Tirole (1992). The principal-agent relationship with an informed principal, II: Common values. *Econometrica 60*(1), 1–42.

Nakamura, E. and J. Steinsson (2005). Price setting in a forward-looking customer markets. *mimeo, Columbia University.*

Okun, A. M. (1981). *Prices and Quantities: A Macroeconomic Analysis.* Brookings Institution.

Peltzman, S. (2000). Prices rise faster than they fall. *Journal of Political Economy 108*(3), 466–502.

Reis, R. (2006). Inattentive producers. *Review of Economic Studies 73*, 793–821.

Reis, R. (2009). A sticky-information general-equilibrium model for policy analysis. K. Schmidt-Heubel and C. Walsh (Eds.). Monetary Policy under Uncertainty and Learning, Central Bank of Chile: forthcoming.

Rotemberg, J. J. (2005). Customer anger at price increases, changes in the frequency of prices changes and monetary policy. *Journal of Monetary Economics* (52), 829–852.

Rotemberg, J. J. (2007). Behavioral aspects of price setting, and their policy implications. *mimeo, Harvard University.*

Rotemberg, J. J. (2008). Fair pricing. *mimeo, Harvard University.*

Wolinsky, A. (1983). Prices as signals of product quality. *Review of Economic Studies, 50*(4), 647–658.

Woodford, M. (2002). Imperfect common knowledge and the effects of monetary policy. P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford (Eds.). Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps, Princeton Univ. Press, 2002.

Woodford, M. (2009). Information-constrained state-dependent pricing. *mimeo, Columbia University.*