

Information, Expectations and the Business Cycle

Rhys Mendes^{*†}

Research Department, Bank of Canada[‡]

PRELIMINARY AND INCOMPLETE – PLEASE DO NOT CITE

December 27, 2007

Abstract

We investigate the ability of a noisy rational expectations model to generate plausible macroeconomic dynamics. The model allows for imperfect, heterogeneous information, and signal extraction from endogenous variables consistent with the assumed trading environment. We find that imperfect information significantly improves the model's ability to generate persistent, hump-shaped responses to a transitory monetary policy shock. This is achieved without the need for mechanical frictions. In addition, the model generates realistic inflation forecast errors.

1. Introduction

The idea that dynamic signal extraction and forecast errors can induce interesting dynamics in macroeconomic models has a long history. Pigou (1929) envisioned

*E-mail: rmendes@bankofcanada.ca

†I am grateful to my advisor, Michelle Alexopoulos, for her patience and guidance. I am also grateful to Michel Juillard and Tom Sargent for helpful suggestions. Amy Corbett provided excellent research assistance for this paper.

‡The views expressed herein are those of the author. No responsibility for them should be attributed to the Bank of Canada.

an economy in which agents' current decisions are based on their forecasts of future variables. He described how shocks could lead to persistent endogenous dynamics through the interaction of forecast errors with production lags. The rational expectations revolution of the 1970s temporarily revived interest in this type of theory. Lucas (1972) used the "island" parable of Phelps (1969) to introduce short-run monetary non-neutrality in a world with rational expectations and flexible prices. In the Phelps-Lucas economy, economic agents cannot directly distinguish between shocks to the relative level of demand on their "island" and aggregate monetary shocks. Instead, they base their decisions on beliefs formed by solving an optimal signal extraction problem. Lucas (1975) further extended things by introducing a model in which misperceptions could persist for more than one period. However, concerns about tractability led Lucas to assume that agents pooled their information at the end of each period. This simplification prevented heterogeneous information from playing an important role in the model's dynamics. Townsend (1983) tackled the heterogeneous information problem head on. Townsend made two contributions. Firstly, he demonstrated that heterogeneous information could introduce interesting economic dynamics. Perhaps more importantly, he showed that in a model with symmetrically dispersed information one quickly runs into an "infinite regress" problem. Agents must forecast the forecasts of others, and forecast the forecasts of the forecasts of others, and so on. The difficulty of formulating and solving such models caused interest in equilibrium signal extraction to wane. For two decades the contributions in this branch of the literature were largely technical: solving the same model as in Townsend by different means.¹

Recent years have witnessed a revival of interest in heterogeneous agent models. The present paper contributes to this nascent resurgence by demonstrating that a model with imperfect and heterogeneous information can generate plausible macroeconomic dynamics without any mechanical frictions. The literature on Dynamic Stochastic General Equilibrium (DSGE) models has emphasized the need for a variety mechanical frictions in order to get these full information rational expectations models to fit the data (e.g., Christiano, Eichenbaum and Evans, 2005; Smets and Wouters, 2003). These mechanical frictions include Calvo-style sticky prices and wages, capital adjustment costs, investment adjustment costs, habit persistence, and dynamic wage and price indexation. In contrast, we study a largely frictionless Phelps-Lucas islands model with monopolistic competition in the goods market. The only mechanical friction is that agents choose consump-

¹See, for example, Sargent (1991) and Kasa (2000).

tion and nominal prices a period in advance. Agents are subject to idiosyncratic technology shocks and relative demand shocks. A given agent trades with a random subset of the other agents in the economy. Agents' information sets include the prices and quantities involved in their own trades; they do not observe prices and quantities for trades to which they are not a party. Agents also observe a set of aggregate statistics (national accounts variables). These aggregate statistics imperfectly measure the true underlying aggregate prices and quantities. With this information set agents are unable to infer the aggregate level of technology with certainty. This leads to inertial behaviour, including endogenously sluggish adjustment of nominal prices, and persistent responses of real variables to shocks. Moreover, inflation expectations exhibit persistent underreaction to shocks (relative to the full information rational expectations case) and autocorrelated forecast errors, consistent with the survey evidence.

The renewed interest in models with imperfect and heterogeneous information began with Woodford's (2002) quantitative demonstration of the potential role of higher-order expectations in generating persistent and hump-shaped responses to monetary shocks. Woodford argues that, even if information about monetary shocks is readily available with only a short delay, imperfect common knowledge can persist due to rational inattention of the type discussed by Sims (2003). A number of authors have subsequently used environments with imperfect common knowledge to study various aspects of macroeconomic dynamics, including Moscarini (2004), Nimark (2005), Bacchetta and van Wincoop (2005), Lorenzoni (2006), Luo (2006), and Mackowiak and Wiederholt (2006).² At the same time, a related literature on sticky information models has developed following Mankiw and Reis (2002). This paper extends the literature by studying a noisy rational expectations model in which *all* agents are fully rational but have imperfect knowledge of the state, and monetary policy is modeled in a fairly standard way.

The model presented in this paper is most closely related to Lorenzoni (2006). In particular, we make extensive use of Lorenzoni's random consumption basket matching technology. We extend Lorenzoni's approach to allow for persistent biases in the baskets. We also increase the set of assets available by introducing nominal bonds. The introduction of nominal bonds allows us to characterize monetary policy by a simple interest rate reaction function. The interest rate is perfectly observed by private agents, but private information on the part of the

²Optimal monetary policy under imperfect common knowledge has been studied by Hellwig (2005), Adam (2007) and Lorenzoni (2007). Amato and Shin (2003) study optimal central bank transparency in a model with heterogeneously informed firms.

central bank creates confusion about the determinants of the interest rate.

The remainder of the paper is organized as follows. Section 2 presents the model. The calibration is discussed in Section 3. Section 4 provides an overview of the solution strategy. Section 5 evaluates the quantitative plausibility of the model by analyzing impulse response functions and forecast errors. Concluding remarks are presented in Section 6.

2. Model

The economy is populated by a continuum of yeoman-farmer households indexed by $i \in [0, 1]$. Each household inhabits a separate island. Each household (island) is a monopolistic supplier of a particular variety of good. Each period, households choose consumption, bonds and the nominal price of their output. Households must choose both the quantities of individual varieties consumed as well as the level of a Dixit-Stiglitz composite of these varieties. We assume that composite consumption and nominal prices are chosen a period in advance. Households must then meet demand for their goods at the chosen price.

At the beginning of each period households observe their idiosyncratic level of technology, the prices they face in the market, the nominal rate of interest, and a set of noisy aggregate statistics. They then choose their current consumption of different varieties given their predetermined level of composite consumption, and their current holdings of bonds and money. They also choose next period's level of composite consumption and nominal prices. At the end of each period, households observe the level of demand for their goods at the posted prices. The central bank's information set includes the aggregate statistics and a private signal on the level of aggregate technology. As will become clear below, this information structure makes it impossible for any agent, including the central bank, to infer the level of aggregate technology with certainty. We also assume that in period t *all* information from period $t - T$ is revealed. This assumption, together with log-linearization of the equilibrium conditions, renders the joint inference-equilibrium problem numerically tractable.

2.1. Households

There is a continuum of households indexed by $i \in [0, 1]$. Households choose consumption, bonds and their nominal wage to maximize:

$$E_{i0} \sum_{t=0}^{\infty} \beta^t \left\{ \log C_{it} + \frac{1}{V_{it}} \log \left(\frac{M_{it}}{\bar{P}_{it}} \right) - \frac{1}{1+\eta} N_{it}^{1+\eta} \right\} \quad (1)$$

where N_{it} is the labour supplied in period t , M_{it} is the stock of money held by household i , \bar{P}_{it} is the average level of consumption prices faced by household i , C_{it} is a consumption composite, and V_{it} is a preference (money demand) shock. The conditional expectation of household i at time t is $E_{it}[\bullet] \equiv E[\bullet | \Omega_{it}]$, where Ω_{it} is household i 's time t information set.

Let lowercase letters denote log deviations from steady state (e.g., $x_{it} = \log(X_{it}/X)$). Then we can write the process for the preference shock as follows:

$$v_{it} = v_t + u_{it}^v \quad (2)$$

$$v_t = \rho_v v_{t-1} + \varepsilon_t^v \quad \varepsilon_t^v \sim N(0, \sigma_v^2) \quad (3)$$

$$u_{it}^v = \rho_{u,v} u_{it-1}^v + \varepsilon_{it}^{u,v} \quad \varepsilon_{it}^{u,v} \sim N(0, \sigma_{u,v}^2) \quad (4)$$

Thus, the preference shock is composed of an aggregate component and an idiosyncratic component. We assume that, at the beginning of each period, household i observes the sum of the two components, v_{it} , but does not directly observe the individual components. The aggregate component, v_t , is essentially a velocity shock. Since v_t is unobserved, the household is left uncertain about aggregate money demand. This, in turn, ensures that observing the aggregate money stock (in conjunction with other observations) will not fully reveal the aggregate state.³

The consumption composite is a Dixit-Stiglitz aggregate of the form:

$$C_{it} = \left(m^{-\frac{1}{\epsilon}} \int_{J_{it-1}} C_{ijt}^{\frac{\epsilon-1}{\epsilon}} dj \right)^{\frac{\epsilon}{\epsilon-1}} \quad (5)$$

³We add the qualifier "in conjunction with other observations" because observing the money stock alone, even in the absence of velocity shocks, will generally not be sufficient to reveal the aggregate state. The intuition is straightforward. In general, if there are N shocks that cause variation in an agents' observed variables, then the agent will require at least N linearly independent observations in order to perfectly determine the N shocks. We setup our model to ensure that we have as many non-technology sources of variation as we have observed variables. Thus, including the technology shock, there are too many shocks for the agents' data to be fully revealing.

where C_{ijt} is the quantity of good j consumed by household i in period t , $J_{it-1} \subset [0, 1]$ is the set of islands from which household i purchases consumption goods in period t (described in more detail below), m is the measure of J_{it-1} , and ϵ is a parameter greater than unity.

Household i produces good i using only its own labour:

$$Y_{it} = A_{it}N_{it}^{1-\alpha} \quad (6)$$

where Y_{it} is output of good i and A_{it} is an island-specific technology factor. Island i 's idiosyncratic level of technology, a_{it} , is given by:

$$a_{it} = a_t + u_{it}^a \quad (7)$$

$$a_t = \rho_a a_{t-1} + \varepsilon_t^a \quad \varepsilon_t^a \sim N(0, \sigma_a^2) \quad (8)$$

$$u_{it}^a = \rho_{u,a} u_{it-1}^a + \varepsilon_{it}^{u,a} \quad \varepsilon_{it}^{u,a} \sim N(0, \sigma_{u,a}^2) \quad (9)$$

Thus, each island's level of technology consists of an aggregate and an idiosyncratic component. Each component is allowed to be persistent. Households are assumed to observe the sum of these two components, a_{it} , at the beginning of each period. They do not observe the two components separately and therefore must form a belief about the aggregate level of technology conditional on their information sets. That a producer knows his own level of technology, but does not directly observe the aggregate, seems no less plausible than the standard full information assumption.

We can now provide a more detailed description of the evolution of the random consumption baskets. Recall that prices are set a period in advance. Thus, if a household observed the current posted prices on every island in the economy, it would be able to learn the aggregate level of technology last period. As in Lorenzoni (2006), we assume that household i observes only a subset of all the consumption prices in the economy – household i only observes the prices at the islands in the random consumption basket, J_{it-1} . In particular, we assume that the average level of time $t - 1$ technology across islands in J_{it-1} is not equal to the mean across all islands in the economy. This assumption ensures that consumption prices (in conjunction with other observations) do not fully reveal the aggregate level of time $t - 1$ technology to any agent.

As mentioned above, the elements of J_{it-1} represent the islands from which household i will purchase consumption goods in period t . We will say that household i has a "consumer relationship" with these islands. At the end of each

period, a consumer relationship is terminated with probability θ (the exogenous separation rate). Consumer relationships that are not terminated continue into the next period. Thus, the measure of consumer relationships that continue from one period to the next is $(1 - \theta)m$. At the beginning of each period, household i receives a set of new consumer relationships, \tilde{J}_{it-1} , of measure θm (this ensures that J_{it-1} is always of measure m). The set of new consumer relationships, \tilde{J}_{it-1} , is selected so that the distribution of $\{\varepsilon_{jt-1}^{u,a} | j \in \tilde{J}_{it-1}\}$ is $N(\zeta_{it-1}, \sigma_{u,a|\zeta}^2)$, and ζ_{it-1} is drawn from a normal distribution.⁴ We refer to ζ_{it} as the "sampling shock". It determines the innovation in the bias of the average productivity level that can be inferred by household i from the prices he directly observes. Thus, the average level of idiosyncratic productivity across all the islands in J_{it} evolves according to:

$$\bar{\varepsilon}_{it}^{u,a} = (1 - \theta) \rho_{u,a} \bar{\varepsilon}_{it-1}^{u,a} + \theta \zeta_{it} \quad (10)$$

where $\rho_{u,a}$ is the autoregressive root on the idiosyncratic component of technology. This bias prevents households from learning the lagged aggregate state through observations of goods prices at the islands they purchase from. In Lorenzoni's (2006, 2007) formulation of the matching technology there was no persistence in consumption baskets ($\theta = 1, \rho_{u,a} = 0$) and therefore the bias was i.i.d. Under our assumptions, consumer relationships last an average of $1/\theta$ periods.⁵ This persistence of consumer relationships makes the bias in the consumption baskets persistent. In addition to being more realistic, the persistent bias makes the household's inference problem more difficult than in the i.i.d. case, thereby slowing learning about the aggregate state.⁶

The household's budget constraint can be written as:

$$B_{it} + M_{it} + \int_{J_{it-1}} P_{jt} C_{ijt} dj = R_{t-1} B_{it-1} + M_{it-1} + P_{it} Y_{it} + TR_t + Q_{it} \quad (11)$$

where R_t is the nominal interest rate between periods $t-1$ and t , B_{it-1} is the bond holdings of household i between periods $t-1$ and t , P_{jt} is the price of good j , TR_t is

⁴The sampling shock for each household, ζ_{it} , is drawn from $N(0, \sigma_\zeta^2)$. The variances must satisfy:

$$\sigma_{u,a}^2 = \sigma_\zeta^2 + \sigma_{u,a|\zeta}^2$$

This restriction ensures that the cross-sectional distribution of technology innovations observed by households is the same as the actual distribution across firms.

⁵Some consumer relationships may be terminated in period t and resumed in period $t+1$. The average length of $1/\theta$ treats these as new relationships.

⁶Learning is faster in the i.i.d. case because the household is receiving many independent observations.

monetary transfers from the central bank, and Q_{it} is a lump-sum state-contingent transfer from the government. Recall that, in period t , all the information from period $t - T$ is revealed. We assume that Q_{it} is contingent upon the revealed state of the economy at time $t - T$ (including idiosyncratic shocks). This transfer is such that it undoes any wealth heterogeneity due to shocks dated $t - T$ and earlier. This has two effects: (1) it eliminates the possibility of individuals violating their "natural borrowing constraint", and (2) it implies that the individual decision rules will be independent of the individual predetermined variable, B_{it-1} . The latter is equivalent to saying that, in the linearized model, individuals will base their consumption decisions solely on their beliefs about the aggregate state of the economy.

The first-order conditions of the problem of minimizing the cost of consuming a given level of C_{it} yield the demand function:

$$C_{ijt} = \left(\frac{1}{m}\right) \left(\frac{P_{jt}}{\bar{P}_{it}}\right)^{-\epsilon} C_{it} \quad \forall j \in J_{it-1} \quad (12)$$

This demand function, along with (5), can be used to derive the cost minimizing price index over the elements of J_{it-1} :

$$\bar{P}_{it} = \left(\frac{1}{m} \int_{J_{it-1}} P_{jt}^{1-\epsilon} dj\right)^{\frac{1}{1-\epsilon}} \quad (13)$$

We can also derive an expression for the total demand for good i . Household i is a monopolistic producer of good i and faces a demand function of the form:

$$\begin{aligned} Y_{it} &= \int_{J_{it}^y} C_{jit} dj \\ &= \int_{J_{it}^y} \left(\frac{1}{m}\right) \left(\frac{P_{it}}{\bar{P}_{jt}}\right)^{-\epsilon} C_{jt} dj \end{aligned}$$

where J_{it}^y is the set of households that consume good i in period t . Assume that the measure of J_{it}^y is subject to shocks, d_{it} , such that we can write:

$$Y_{it} = P_{it}^{-\epsilon} e^{d_{it}} \int_0^1 \bar{P}_{jt}^{\epsilon} C_{jt} dj \quad (14)$$

where,

$$d_{it} = \rho_d d_{it-1} + \varepsilon_{it}^d \quad \varepsilon_{it}^d \sim N(0, \sigma_d^2)$$

Thus, d_{it} acts as a relative demand shock for good i . As in Lorenzoni (2006), this shock prevents the household's sales from revealing the lagged aggregate level of technology.

Recall that consumption and nominal prices are chosen a period in advance. The first-order conditions for C_{it} and P_{it} are:

$$\frac{1}{C_{it}} = E_{it-1} \left[\widehat{\Phi}_{it} \right] \quad (15)$$

$$E_{it-1} \left[\widehat{\Phi}_{it} Y_{it} \left(\frac{P_{it}}{\bar{P}_{it}} - \left(\frac{\epsilon}{\epsilon - 1} \right) \left(\frac{1}{1 - \alpha} \right) \frac{(Y_{it}/A_{it})^{\frac{\eta+\alpha}{1-\alpha}}}{A_{it} \widehat{\Phi}_{it}} \right) \right] = 0 \quad (16)$$

where,

$$\widehat{\Phi}_{it} \equiv \Phi_{it} \bar{P}_{it}$$

and Φ_{it} is the Lagrange multiplier on the budget constraint. The first-order conditions for bonds and money imply:

$$\widehat{\Phi}_{it} = \beta E_{it} \left[\widehat{\Phi}_{it+1} \left(\frac{\bar{P}_{it}}{\bar{P}_{it+1}} \right) \right] R_t \quad (17)$$

$$\frac{M_{it}}{\bar{P}_{it}} = \frac{1}{V_{it}} \left(\frac{R_t}{R_t - 1} \right) \frac{1}{\widehat{\Phi}_{it}} \quad (18)$$

As in full information models, the additive separability of money balances in the utility function implies that money will not affect aggregate dynamics. Money is included in the model only because it plays a role as an information variable.

Private sector behaviour can be characterized by equations (11), (13), (14), (15), (16), (17), and (18). The evolution of household information sets will be detailed after the descriptions of the monetary authority and the statistical agency. Market clearing in the money and bond markets implies:

$$\begin{aligned} M_t &= \int_0^1 M_{it} di \\ 0 &= \int_0^1 B_{it} di \end{aligned}$$

where M_t is the aggregate quantity of money supplied by the central bank. It will also be useful to define some aggregates. For our purposes, we need only

define log-linear approximations to the aggregates:

$$\begin{aligned} y_t &= \int_0^1 y_{it} di \\ p_t &= \int_0^1 p_{it} di \\ \pi_t &= p_t - p_{t-1} \end{aligned}$$

These aggregates will be used in the descriptions of the monetary authority, the statistical agency and the information sets of different agents.

2.2. Monetary Authority

The 1970s vintage noisy rational expectations models of the monetary transmission mechanism, such as Lucas (1972), fell out of favor in part because they were unable to generate persistent effects of monetary shocks. This failure was due to the fact that, in these models, agents' observation of the policy instrument allowed them to infer the true state of the economy. In the 1970s vintage models, the money supply was usually the policy instrument. Since, in practice, the money supply is observed with a very short lag, these models implied virtually no persistence.⁷ In modern treatments of monetary policy, a short-term nominal interest rate is usually the policy instrument. The short-term rate is observed with no lag at all, thus one might expect this to close the door to imperfect information models of the transmission mechanism. However, in our setup, neither the nominal interest rate nor the money stock fully reveal the aggregate state. We assume that the central bank observes a private noisy signal on the level of aggregate technology. This assumption, combined with our specification of the monetary policy reaction function, ensures that the nominal interest rate is not fully revealing (in conjunction with other observations).⁸

The central bank sets the nominal interest rate according to (in linearized form):

$$r_t = \rho_r r_{t-1} + (1 - \rho_r) \left(E_t^{cb} [r_t^n + \pi_{t+1}^*] + \gamma_\pi (E_t^{cb} [\pi_{t+1}] - \pi_t^*) + \gamma_x E_t^{cb} [x_t] \right) + \varepsilon_t^r \quad (19)$$

⁷As Mankiw, Runkle and Shapiro (1984) note, the initial money supply data is subject to revision, which suggests that it is error-ridden. This would make it less than fully revealing. We do not take up this line of argument here.

⁸The velocity shocks discussed in the previous section ensure that the money stock is not fully revealing.

where $E_t^{cb}(\bullet)$ is the expectation of the central bank conditional on its time t information set, π_t^* is the inflation target, x_t is the output gap, and ε_t^r is a transitory monetary policy shock. The output gap is defined as $x_t \equiv y_t - y_t^n$, where y_t^n is the natural rate of output that would prevail under full information with no decision delays. The two policy shocks evolve according to:

$$\pi_t^* = \rho_\pi \pi_{t-1}^* + \varepsilon_t^\pi \quad \varepsilon_t^\pi \sim N(0, \sigma_{\pi^*}^2) \quad (20)$$

$$\varepsilon_t^r \sim N(0, \sigma_r^2) \quad (21)$$

Naturally, the central bank is assumed to know the values of both policy shocks at each point in time. However, private agents must form inferences on these shocks based on their observations of central bank behaviour. The central bank also observes the aggregate statistics described in the next section and a private signal on the level of aggregate technology, s_t^{cb} :

$$s_t^{cb} = a_t + u_t^{cb} \quad (22)$$

$$u_t^{cb} = \rho_{u,cb} u_{t-1}^{cb} + \varepsilon_t^{u,cb} \quad \varepsilon_t^{u,cb} \sim N(0, \sigma_{u,cb}^2) \quad (23)$$

This private signal prevents the nominal interest rate from completely revealing the sum of the policy shocks. When households/firms observe an unexplained monetary policy movement, they cannot determine with certainty whether the cause was a policy shock or the response of the central bank to its private signal.

2.3. Statistical Agency and Information Sets

We assume that there exists a statistical agency that publishes information on aggregate variables. The agency's first release of data does not correspond exactly to the true values of the economic variables being measured. It is assumed that data that reflects the true values becomes available after T periods. Given that revisions to early national accounts data can be quite substantial, we believe this is a realistic assumption. We must make a choice about how to model the deviation of the statistics from the true values.

Mankiw, Runkle and Shapiro (1984) and Mankiw and Shapiro (1986) propose two extreme characterizations of the revision process: *news* and *noise*. Under the news interpretation, the statistical agency uses its information to form an optimal forecast of the true value. Revisions to preliminary data reflect news that arrives after the data is published. In contrast, under the noise interpretation, preliminary data are contaminated with measurement error that is uncorrelated

with the true value. Thus, under this view, the preliminary data are not optimal estimates of the true values and revisions reduce the measurement error by drawing on larger samples, correcting mistakes, etc.

One can examine the validity of these characterizations by using forecast efficiency tests.⁹ In particular, one can test the hypothesis that revisions are uncorrelated with the preliminary data (as they should be under the news characterization). In addition, one can examine the variance of different vintages of data. If the preliminary data are an optimal forecast then the variance of the revised data should be greater than the variance of the preliminary data. Using data from 1975 to 1982, Mankiw and Shapiro conclude that US GNP revisions are better characterized as news rather than noise. Faust, Rogers and Wright (2005) use a longer dataset to confirm the Mankiw-Shapiro results for the US. However, they also find that revisions in other G-7 countries might be better modeled as noise. Croushore and Stark (2001) analyze revisions to US consumption data using a 30 year real-time data set. They find that revisions up to one year after the initial release are uncorrelated with the preliminary data, while subsequent revisions are weakly correlated with the preliminary data. They also find that the standard deviation of consumption growth rises from the initial release to one year revision, but it falls in all subsequently revised data. Thus, their results are mixed but point more toward the noise interpretation for consumption revisions.

In our simple model output is equal to consumption. Hence, we put more weight on the Croushore-Stark results and model the revisions as being due to measurement error. In period t the statistical agency publishes statistics of the form:

$$\tilde{g}_{t-1,t}^y = y_{t-1} - y_{t-2} + \tilde{u}_{1t}^y \quad (24)$$

$$\tilde{\pi}_{t-1,t} = \pi_{t-1} + \tilde{u}_{1t}^\pi \quad (25)$$

where,

$$\tilde{u}_{1t}^y = \rho_{\tilde{y},1} \tilde{u}_{1t-1}^y + \tilde{\varepsilon}_{1t}^y \quad \tilde{\varepsilon}_{1t}^y \sim N(0, \sigma_{u,\tilde{y}}^2) \quad (26)$$

$$\tilde{u}_{1t}^\pi = \rho_{\tilde{\pi},1} \tilde{u}_{1t-1}^\pi + \tilde{\varepsilon}_{1t}^\pi \quad \tilde{\varepsilon}_{1t}^\pi \sim N(0, \sigma_{u,\tilde{\pi}}^2) \quad (27)$$

Both individuals and the central bank observe these statistics in each period. The statistics give agents a noisy measurement of each series.

⁹See, e.g., Mincer and Zarnowitz (1969).

2.4. Information Sets

We can now return to the definition of the household and central bank information sets. At the beginning of each period, household i observes the exogenous variables a_{it} and v_{it} . It also observes the nominal interest rate, r_t , consumption prices, $\{p_{jt}\}_{j \in J_{it-1}}$, the lagged aggregate money stock, m_{t-1} , and the two noisy aggregate statistics, $\tilde{g}_{t-1,t}^y$ and $\tilde{\pi}_{t-1,t}$. At the end of each period household i observes its sales, y_{it} . It is easy to verify that \bar{p}_{it} is a sufficient statistic for $\{p_{jt}\}_{j \in J_{it-1}}$, so we can write household i 's information set at the beginning of period t as:

$$\Omega_{it} = \{a_{it}, v_{it}, r_t, \bar{p}_{it}, m_{t-1}, \tilde{g}_{t-1,t}^y, \tilde{\pi}_{t-1,t}, y_{it-1}\} \cup \Omega_{it-1} \cup \mathfrak{S}_{t-T} \quad (28)$$

where \mathfrak{S}_{t-T} is the set of all innovations dated $t - T$ and earlier. The inclusion of \mathfrak{S}_{t-T} captures our assumption that all information is revealed after T periods. Note that we assume that the lagged money stock is observed without error. This assumption is manifestly false, but adding measurement error makes little difference.¹⁰

The information set of the central bank was described in the previous subsection and can be summarized as follows:

$$\Omega_t^{cb} = \{\pi_t^*, \varepsilon_t^r, s_t^{cb}, m_{t-1}, \tilde{g}_{t-1,t}^y, \tilde{\pi}_{t-1,t}\} \cup \Omega_{t-1}^{cb} \cup \mathfrak{S}_{t-T} \quad (29)$$

Note that r_t is not included in the central bank's information set because it is not conditioning information for the central bank – it is a choice variable. The current and lagged nominal interest rates are functions of Ω_t^{cb} and therefore need not be included.

3. Calibration

The model period is a quarter. The discount factor, β , is set to 0.9925 in order to generate a steady-state real interest rate of 3 per cent. The preference parameter, η , is set to 0.33 to obtain a Frisch labour supply of 3.

The production parameter, α , is chosen to imply a labour share of 2/3 (given the value of ϵ discussed below). The calibration of the stochastic processes for the idiosyncratic and aggregate technology shocks follows Kahn and Thomas (2006). The parameters ρ_a and σ_a are estimated by computing Solow residuals using

¹⁰The reason for this is that the velocity shocks act like measurement error on the money stock.

NIPA data on US real GDP and private capital, along with the total employment hours series constructed by Prescott, Ueberfeldt and Cociuba (2005) from CPS household survey data. This yields $\rho_a = 0.859$ and $\sigma_a = 0.014$. As Kahn and Thomas (2006) point out, there is little agreement on the correct persistence for the idiosyncratic technology shock. They simply set the persistence of the idiosyncratic shock equal to the persistence of the aggregate shock. We do the same. We also set the standard deviation of the innovation to the idiosyncratic shock process, $\sigma_{u,a}$, equal to 0.022 as in Kahn and Thomas.

The parameter $\epsilon = 11$ is chosen to generate a 10 per cent markup. In order to calibrate the stochastic process for the relative demand shocks, we adapt the calibration used by Busato (2004). Busato examined relative demand shocks to two categories of consumption: Food and Clothing/Shoes. We use the midpoint of his estimates for the two sectors to calibrate our relative demand shock process. This yields $\rho_d = 0.96$ and $\sigma_d = 0.00712$. These numbers are conservative in the sense that we would expect relative demand shocks at the firm level to be more volatile than at the level of broad consumption categories (because the product of any given firm is likely to have many close substitutes, while food is not a good substitute for shoes).

To choose a value for T , we consider the revision process in the US. The Bureau of Economic Analysis (BEA) releases annual revisions to GDP data for 3 years after the initial release. After 5 years the BEA releases a comprehensive revision that is benchmarked to input-output tables. We assume that the 5 year revisions reflect true values (free of measurement error).¹¹ For this reason, we set $T = 20$. That is, at time t the true state of the economy at $t - 20$ is fully revealed. Faust, Rogers and Wright (2005) find that the root mean squared revision between final and preliminary data for US quarterly GDP growth is 0.53 per cent. Thus, we set $\sigma_{u,\tilde{y}} = 0.0053$ and $\rho_{\tilde{y},1} = 0$. Croushore (2007) analyzes revisions to measures of PCE inflation. Based on this analysis, we set $\sigma_{u,\tilde{\pi}} = 0.005$ and $\rho_{\tilde{\pi},1} = 0$.

We set the parameters of the policy rule to be consistent with the post-1979 estimates reported by Clarida, Gali and Gertler (1999). In particular, we set $\gamma_\pi = 1.5$, $\gamma_x = 0.1$ and $\rho_r = 0.85$. Kozicki and Tinsley (2005) estimate the variances of the inflation target and the transitory policy shock in an empirical model with asymmetric information. Using the estimates in Kozicki and Tinsley

¹¹Of course, we know that this is not true. If the 5 year revisions were in fact free of measurement error then the statistical discrepancy should be zero. It is not. Our results regarding the importance of state uncertainty would only be strengthened by acknowledging that the true state is never revealed.

(2005), we set $\sigma_{\pi^*} = 0.0021/4$ and $\sigma_r = 0.0076/4$. They assume the target shock is permanent. We approximate this assumption by setting $\rho_\pi = 0.95$. We do not have any information that would help to calibrate the parameters of the error process in the central bank's private signal. Assuming that the central bank's private signal is as informative as private agents' observation of their idiosyncratic level of technology seems to be a natural starting point. Hence, we set $\rho_{u,cb} = \rho_{u,a}$ and $\sigma_{u,cb} = \sigma_{u,a}$.

Kleshchelski and Vincent (2007) calibrate their market share model so that consumer relationships last an average of 3.5 years. For this reason, we set the exogenous separation rate for consumer relationships, θ , to 0.07. To calibrate the standard deviation of the sampling shock, σ_ζ , we look at the evidence on cross-sectional variance of inflation rates reported by Hobijn and Lagakos (2005). They compute household-specific inflation rates, which we take as being equivalent to the model object, $\bar{p}_{it} - \bar{p}_{it-1}$. Taking as given all the other parameter values, we choose σ_ζ to generate a cross-sectional variance of household inflation rates of 0.7 percentage points, as in Hobijn and Lagakos.

This calibration of the model obviously relies on some rough approximations and back-of-the-envelope calculations. However, we have taken the calibration more seriously than has been standard in the literature on imperfect information macroeconomic models. For this reason, we regard our analysis as a quantitatively plausible, though tentative, indication of the promise of this modelling approach.

4. Solution Method

This section provides a brief overview of the numerical procedure used to solve the model. A technical appendix with further details is available from the author upon request. The solution method relies on log-linearization of the equilibrium conditions and the truncation method of Townsend (1983). We assume that after T periods the true state of the economy is revealed. That is, at time t all agents learn all of the shocks from time $t - T$ and earlier. This circumvents the problem of forecasting the forecasts of others identified by Townsend. One can then apply standard methods for solving agents' inference problems, and the equilibrium can be computed using a method of undetermined coefficients.

We begin by log-linearizing the equilibrium conditions. In the log-linear approximation, the assumed normality of the idiosyncratic shocks then implies that cross-sectional distributions are normal with constant variance and time-varying mean. This, in turn, implies that we need only track the mean of the distribution

over time – significantly simplifying the equilibrium computation problem.

In addition to assuming that all information is revealed after T periods, we also assume that the impulse response functions flatten out after $\bar{T} = T + T_2$ periods. Then we can conjecture linear decision rules of the form:

$$p_{it+1} = \varphi_{u,1}^p E_{it} [\mathbf{u}_{it}^{[T]}] + \varphi_{u,2}^p \mathbf{u}_{it-T}^{[T_2]} + \varphi_{U,1}^p E_{it} [\mathbf{U}_t^{[T]}] + \varphi_{U,2}^p \mathbf{U}_{t-T}^{[T_2]} + \bar{\psi}_p \mathbf{X}_{t-\bar{T}} \quad (30)$$

$$c_{it+1} = \varphi_{u,1}^c E_{it} [\mathbf{u}_{it}^{[T]}] + \varphi_{u,2}^c \mathbf{u}_{it-T}^{[T_2]} + \varphi_{U,1}^c E_{it} [\mathbf{U}_t^{[T]}] + \varphi_{U,2}^c \mathbf{U}_{t-T}^{[T_2]} + \bar{\psi}_c \mathbf{X}_{t-\bar{T}} \quad (31)$$

$$\widehat{\phi}_{it} = \varphi_{u,1}^\phi E_{it} [\mathbf{u}_{it}^{[T]}] + \varphi_{u,2}^\phi \mathbf{u}_{it-T}^{[T_2]} + \varphi_{U,1}^\phi E_{it} [\mathbf{U}_t^{[T]}] + \varphi_{U,2}^\phi \mathbf{U}_{t-T}^{[T_2]} + \bar{\psi}_\phi \mathbf{X}_{t-\bar{T}} \quad (32)$$

where,

$$\mathbf{U}_t^{[T]} \equiv \begin{bmatrix} \mathbf{U}_t \\ \mathbf{U}_{t-1} \\ \vdots \\ \mathbf{U}_{t-T+2} \\ \mathbf{U}_{t-T+1} \end{bmatrix}, \quad \mathbf{U}_{t-T}^{[T_2]} \equiv \begin{bmatrix} \mathbf{U}_{t-T} \\ \mathbf{U}_{t-T-1} \\ \vdots \\ \mathbf{U}_{t-\bar{T}+2} \\ \mathbf{U}_{t-\bar{T}+1} \end{bmatrix} \quad (33)$$

and $\mathbf{U}_t \equiv (\varepsilon_t^a, \varepsilon_t^v, \varepsilon_t^\pi, \varepsilon_t^r, \varepsilon_t^{u,cb}, \widetilde{\varepsilon}_{1t}^y, \widetilde{\varepsilon}_{1t}^\pi)'$. The stacked vectors $\mathbf{u}_{it}^{[T]}$ and $\mathbf{u}_{it-T}^{[T_2]}$ are defined analogously, with $\mathbf{u}_{it} \equiv (\varepsilon_{it}^{u,a}, \varepsilon_{it}^{u,v}, \varepsilon_{it}^d, \zeta_{it})'$. We also define,

$$\mathbf{X}_{t-\bar{T}} = \mathbf{U}_{t-\bar{T}} + \mathbf{X}_{t-\bar{T}-1}$$

The recursive equation for $\mathbf{X}_{t-\bar{T}}$ captures the long-run impact of aggregate shocks. In our setup, only nominal prices and the price-level will not return to zero in the long-run. We do not include decision rules for individual money and bond holdings because they can be computed residually.

We also conjecture aggregate laws of motion of the form:

$$y_{t+1} = \Phi_y \mathbf{U}_t^{[T_1]} + \Psi_y \mathbf{U}_{t-T_1}^{[T_2]} + \bar{\Psi}_y \mathbf{X}_{t-\bar{T}} \quad (34)$$

$$p_{t+1} = \Phi_p \mathbf{U}_t^{[T_1]} + \Psi_p \mathbf{U}_{t-T_1}^{[T_2]} + \bar{\Psi}_p \mathbf{X}_{t-\bar{T}} \quad (35)$$

$$\widehat{\phi}_t = \Phi_\phi \mathbf{U}_t^{[T_1]} + \Psi_\phi \mathbf{U}_{t-T_1}^{[T_2]} + \bar{\Psi}_\phi \mathbf{X}_{t-\bar{T}} \quad (36)$$

$$r_t = \Phi_r \mathbf{U}_t^{[T_1]} + \Psi_r \mathbf{U}_{t-T_1}^{[T_2]} + \bar{\Psi}_r \mathbf{X}_{t-\bar{T}} \quad (37)$$

We can then use these laws of motion, together with the household first-order conditions, to solve for the coefficients of the decision rules as functions of the

coefficients of the laws of motion. We can then solve the inference problems of households and the central bank conditional on these conjectured laws of motion. That is, we can compute the linear projection of the unobserved shocks on the observed variables. This gives:

$$E_{it} \left[\mathbf{U}_t^{[T_1]} \right] = \Xi_{U,1} \mathbf{u}_{it}^{[T_1]} + \Xi_{U,2} \mathbf{U}_t^{[T_1]} \quad (38)$$

$$E_{it} \left[\mathbf{u}_{it}^{[T_1]} \right] = \Xi_{u,1} \mathbf{u}_{it}^{[T_1]} + \Xi_{u,2} \mathbf{U}_t^{[T_1]} \quad (39)$$

$$E_t^{cb} \left[\mathbf{U}_t^{[T_1]} \right] = \Xi_U^{cb} \mathbf{U}_t^{[T_1]} \quad (40)$$

We can substitute these expressions for the expectations into the decision rules and the monetary policy rule, (19), and aggregate in order to get expressions that are of the same form as the conjectured laws of motion. We can then match coefficients and iterate to convergence.

This procedure can be summarized as follows:

1. Conjecture individual decision rules, conditional on beliefs about shocks
2. Conjecture linear aggregate laws of motion
3. Solve for the coefficients of the individual decision rules as functions of the coefficients of the aggregate laws of motion.
4. Solve the inference problems of all agents conditional on the aggregate laws of motion (i.e., compute the linear projection of unobserved shocks on observed variables)
5. Use solution of inference problems from (4) to eliminate expectations from individual decision rules and the monetary policy rule, and aggregate to obtain updated aggregate laws of motion.
6. Stop if aggregate laws of motion have converged, else return to (4).

5. Results

Figures 1-4 show the responses of inflation, output and the nominal interest rate to a transitory monetary policy shock (ε_t^r), an inflation target shock (ε_t^π), a shock

to the central bank's private signal ($\varepsilon_t^{u,cb}$), and a technology shock (ε_t^a).¹² The first three of these are all shocks to monetary policy. CEE estimate a VAR and identify the impact of a monetary policy shock using restrictions on contemporaneous responses. Their VAR results are broadly representative of the empirical literature. For the purposes of evaluating our model, three aspects of their results are important:

1. the response of output is hump-shaped, peaking after roughly one and a half years at about 0.5 per cent, and returning to control after about three years;
2. the response of inflation is hump-shaped, peaking at roughly 0.2 per cent, two years after the shock;
3. the interest rate falls for roughly one year.

CEE find that, in order to replicate the hump-shaped responses of output and inflation, they must introduce a role for lagged endogenous variables in the structural equations. They do so by assuming habit persistence, investment adjustment costs, and indexation of wages and prices to lagged inflation.

Figure 1 displays our model's predicted responses to a transitory monetary policy shock (ε_t^r) for some key variables. The shock is scaled to 75 basis in order to be directly comparable with CEE's results. The output response peaks at roughly 0.25 in the quarter immediately following the shock. The response of inflation exhibits the hump-shape characteristic of the VAR response, but it peaks too quickly and too high. The model response peaks at 0.33 only 5 quarters after the shock. It is worth noting that, consistent with the VAR evidence, the model generates a temporary decline in inflation after an expansionary monetary policy shock. This phenomenon is often referred to as the "price puzzle." Our model provides a novel explanation for this puzzle. Upon observing a decline in the nominal interest rate, price-setters are uncertain about the cause of the decline. They initially assign greater weight to the possibility that it is due to a positive technology shock, rather than a monetary policy shock. That is, the average price-setter believes he has received a negative idiosyncratic technology shock, but the aggregate level of technology has increased. As time goes on, the price-setters learn that it is more likely a monetary policy shock, and inflation

¹²The horizontal axes are denominated in quarters, while the vertical axes report percentage deviations from steady-state.

rises. Finally, the nominal interest rate falls for about a year and a half – within the confidence intervals reported in CEE.

Looking at the results from the limited information model in isolation obscures the full impact of the informational assumptions. Figures 5-8 plot impulse responses from the limited information model along side responses from the full information analog. The full information analog is constructed by assuming that all date t innovations are perfectly observed by all agents at date t . It is immediately evident that the full information version of the model cannot generate any inflation inertia – all of the inertia in our model is coming from the informational assumptions, not the mechanical features of the model. Moreover, in the full information model output does not respond at all to the policy shock. Clearly, the informational assumptions have a profound impact on model dynamics and move the model’s responses in the right direction.

Furthermore, the model implies very persistent responses of inflation and output to an inflation target shock (Figure 2). In this case, output has a hump-shaped profile and peaks just below 0.6 about a year and a half after the shock. Inflation peaks well after output, about 2.5 years after the shock. The persistent response of output cannot be attributed to any of the mechanical features of the model. In fact, as Figure 6 shows, output does not respond at all in the full information version of the model. Thus, the informational assumptions are completely responsible for the propagation of the inflation target shock to output.

The CEE VAR does not explicitly distinguish between transitory policy shocks and persistent inflation target shocks. It seems likely that, in reality, there have been shifts in the Federal Reserve’s implicit inflation objective over recent decades (see, e.g., Kozicki and Tinsley, 2005). It is not immediately obvious what the implications of ignoring inflation target shocks are for the VAR impulse responses. One way of addressing this question would be to estimate VARs on simulated data generated from our model. We plan to take up this line of inquiry in future versions of this paper.

It is somewhat surprising that our simple model performs as well as it does. Recall that the model has no lagged endogenous variables at all. The inertia evident in the impulse responses results from endogenous inertia in expectations. We can analyze this further by examining the behaviour of inflation forecast errors.

Mankiw, Reis and Wolfers (2006) attempt to ascertain the extent of persistence in actual forecast errors by using survey data on inflation expectations. They use monthly data to estimate an equation of the form:

$$\pi_t - \bar{E}_{t-12}\pi_t = \alpha + \beta (\pi_{t-12} - \bar{E}_{t-24}\pi_{t-12}) \quad (41)$$

where the expectations are median survey expectations. They use four different surveys in the analysis. They find that α is always insignificant and that β is always significant and varies between 0.37 and 0.64 depending on the survey used.

We run a similar regression on simulated quarterly data generated from our model:

$$\pi_t - \bar{E}_{t-4}\pi_t = \alpha + \beta (\pi_{t-4} - \bar{E}_{t-8}\pi_{t-4}) \quad (42)$$

where the time subscripts have been modified because our simulated data is quarterly, not monthly. The regression yields $\beta = 0.19$ for our model. Note that for *any* standard full information DSGE model, we would find $\beta = 0$. This is simply due to the fact that, under full information, forecast errors can be due solely to unanticipated innovations. However, in our limited information model forecast errors depend on current and past innovations and can therefore be autocorrelated. Past innovations affect current forecast errors because agents are continuously learning about current *and* past innovations. Thus, our model is able to generate a more realistic autocorrelation of inflation forecast errors than a standard full information model. Analysis of forecast errors provides a simple and useful way of discriminating between limited and full information models that might be approximately observationally equivalent in terms of their predictions for other variables.

6. Conclusion

We have demonstrated that introducing imperfect and heterogeneous information moves a simple, frictionless model closer to replicating several important business cycle facts. The model is able to achieve this without the need for any mechanical frictions like Calvo pricing, indexation or adjustment costs. In fact, the model's structural equations contain no lagged endogenous variables. All the inertia generated by the model is due to expectational inertia. In addition, the model is also closer to matching the survey evidence on the autocorrelation of the cross-sectional average of inflation forecast errors than standard DSGE models. This suggests that the mechanism by which our model generates inertia might be more realistic than the mechanisms used by full information models.

Other types of models with information imperfections can also generate substantial and realistic persistence. In particular, both Rational Inattention and Sticky Information models hold promise. However, these models introduce a friction to generate expectational inertia. In contrast, our model is based on the

simple assumption that the data available to individuals does not fully reveal the state of the economy. We have found that accounting for this fact can have a substantial impact on macroeconomic dynamics. Before introducing information imperfections of the Rational Inattention/Sticky Information sort, we should fully explore how far we can get on the basis of imperfect data alone. Nevertheless, our model is merely suggestive of what might be possible with models that account for the limited information sets available to individuals in reality. This line of research appears promising, but much more work is required in this direction.

References

1. Adam, K. (2007): Optimal Monetary Policy with Imperfect Common Knowledge, *Journal of Monetary Economics*, 54(2), pp. 276-301.
2. Amato, J.D., and H.S. Shin (2003): Public and Private Information in Monetary Policy Models, BIS Working Paper No. 138.
3. Bacchetta, P. and E. van Wincoop (2005): Can Information Heterogeneity Explain the Exchange Rate Determination Puzzle?, *American Economic Review*, 96, pp. 552-576.
4. Busato, F. (2004): Relative Demand Shocks, University of Aarhus Economics Working Paper No. 2004-11.
5. 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, pp. 1-45.
6. Clarida, R., J. Gali and M. Gertler (2000): Monetary Policy Rules And Macroeconomic Stability: Evidence And Some Theory, *Quarterly Journal of Economics*, Vol. 115(1), pp. 147-180.
7. Croushore, D. and T. Stark (2001): A real-time data set for macroeconomists, *Journal of Econometrics*, vol. 105, Issue 1, pp. 111-130.
8. Croushore, D (2007): Revisions to PCE inflation measures: implications for monetary policy, Unpublished Manuscript, University of Richmond.
9. Faust, J., J.H. Rogers and J.H Wright (2005): News and Noise in G-7 GDP Announcements, *Journal of Money, Credit and Banking*, Vol. 37, pp. 403-419.
10. Hellwig, C. (2005): Heterogeneous Information and the Benefits of Public Information Disclosures, Unpublished Manuscript, UCLA.
11. Hobijn, B. and D. Lagakos (2005): Inflation Inequality in the United States, *Review of Income and Wealth*, 51(4), pp. 581-606.

12. Kahn, A. and J.K. Thomas (2007): Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics, NBER Working Papers 12845.
13. Kasa, K. (2000): Forecasting the Forecasts of Others in the Frequency Domain, *Review of Economic Dynamics*, 3, pp. 726-56.
14. Kleshchelski, I. and N. Vincent (2007): Market Share and Price Rigidity, Unpublished Manuscript, Northwestern University.
15. Kozicki, S. and P. Tinsley (2005): Permanent and transitory policy shocks in an empirical macro model with asymmetric information, *Journal of Economic Dynamics and Control*, 29(11), pp. 1985-2015.
16. Lorenzoni, G. (2006): A Theory of Demand Shocks, Unpublished Manuscript, MIT.
17. Lorenzoni, G. (2007): News Shocks and Optimal Monetary Policy, NBER Working Paper No. 12898.
18. Lucas, R.E., (1972): Expectations and Neutrality of Money, *Journal of Economic Theory*, 4(2), pp. 103-24.
19. Lucas, R.E., (1975): An Equilibrium Model of the Business Cycle, *Journal of Economic Theory*, 4(2), pp. 103-24.
20. Luo, Y. (2006): Consumption Dynamics under Information Processing Constraints, forthcoming in the *Review of Economic Dynamics*.
21. Mackowiak, B. and M. Wiederholt (2006): Optimal Sticky Prices under Rational Inattention, Unpublished Manuscript, Northwestern University
22. Mankiw, N.G., D.E. Runkle and M.D. Shapiro (1984): Are Preliminary Announcements of the Money Stock Rational Forecasts?, *Journal of Monetary Economics*, 14, pp.15-27.
23. Mankiw, N.G. and M.D. Shapiro (1986): News or Noise: An Analysis of GNP Revisions, *Survey of Current Business*, May 1986, pp.20-25.
24. Mankiw, N.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), pp. 1295-1328.

25. Mankiw, N.G. and R. Reis (2006): Pervasive Stickiness, *American Economic Review*, 96(2), pp. 164-169.
26. Mankiw, N.G., R. Reis and J. Wolfers (2006): Disagreement about Inflation Expectations, NBER Working Paper No. 9796.
27. Mincer, J. and V. Zarnowitz (1969): The Evaluation of Economic Forecasts in J. Mincer (ed.), *Economic Forecasts and Expectations*, NBER, New York.
28. Moscarini, G. (2004): Limited Information Capacity as a Source of Inertia, *Journal of Economic Dynamics and Control*, Vol. 28(10), pp. 2003-2035.
29. Nimark, K.P. (2005): Calvo Pricing and Imperfect Common Knowledge: A Forward Looking Model of Rational Inflation Inertia, ECB Working Paper No. 474.
30. Phelps, E.S. (1969): Introduction: The New Microeconomics in Employment and Inflation Theory, in E.S. Phelps et al.: *Microeconomic Foundations of Employment and Inflation Theory*, Norton, New York.
31. Pigou, A.C. (1929): *Industrial Fluctuations*, MacMillan and Co. Ltd., London.
32. Prescott, E.C., A. Ueberfeldt and S. Cociuba (2005): U.S. Hours and Productivity Behavior Using CPS Hours Worked Data: 1959-I to 2005-II, Federal Reserve Bank of Minneapolis Mimeo.
33. Sargent, T.J. (1991) Equilibrium with Signal Extraction from Endogenous Variables, *Journal of Economic Dynamics and Control*, 15(2), pp. 245-273.
34. Sims, C. (2003): Implications of Rational Inattention, *Journal of Monetary Economics*, 50 (3), pp. 665-690.
35. Smets, F. and Wouters, R. (2003): An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area, *Journal of the European Economic Association*, Vol. 1(5), pages 1123-1175.
36. Townsend, R.M. (1983): Forecasting the Forecasts of Others, *Journal of Political Economy*, 91(4), pp. 546-587.

37. Woodford, M. (2002): Imperfect Common Knowledge and the Effects of Monetary Policy, in P. Aghion, R. Frydman, J. Stiglitz and M. Woodford, eds., *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, Princeton: Princeton University Press.

Figure 1: Interest Rate Shock (75 bps)

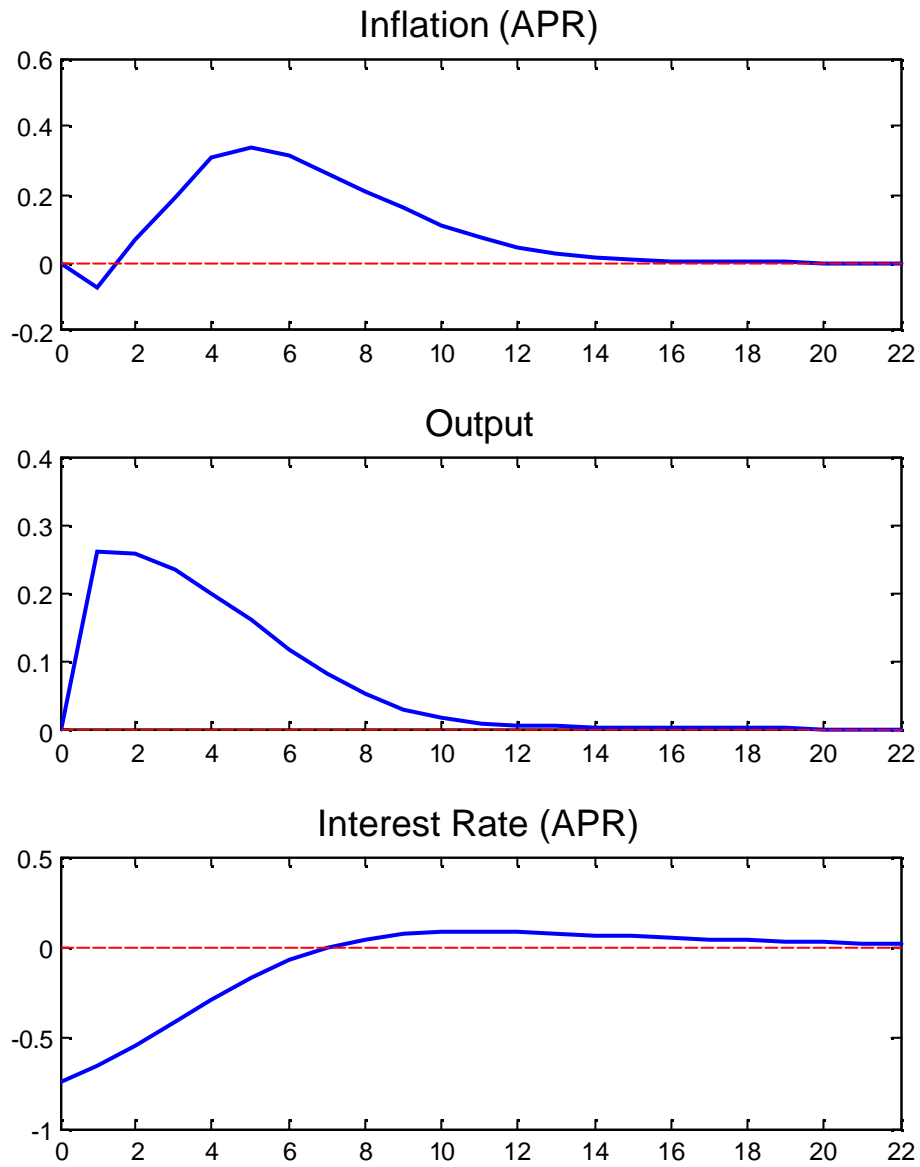


Figure 2: Inflation Target Shock (1 per cent)




Figure 3: Central Bank Signal Shock (1 per cent)

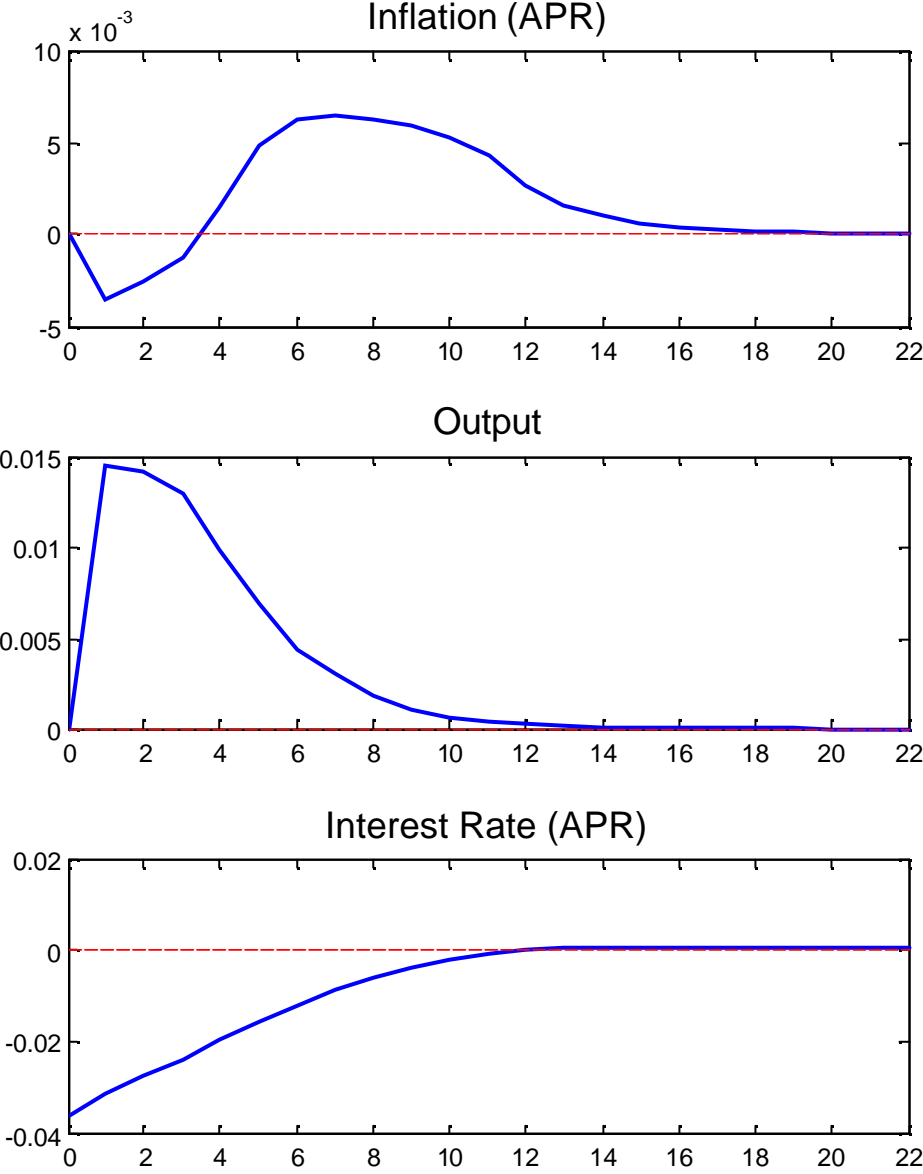


Figure 4: Technology Shock (1 per cent)

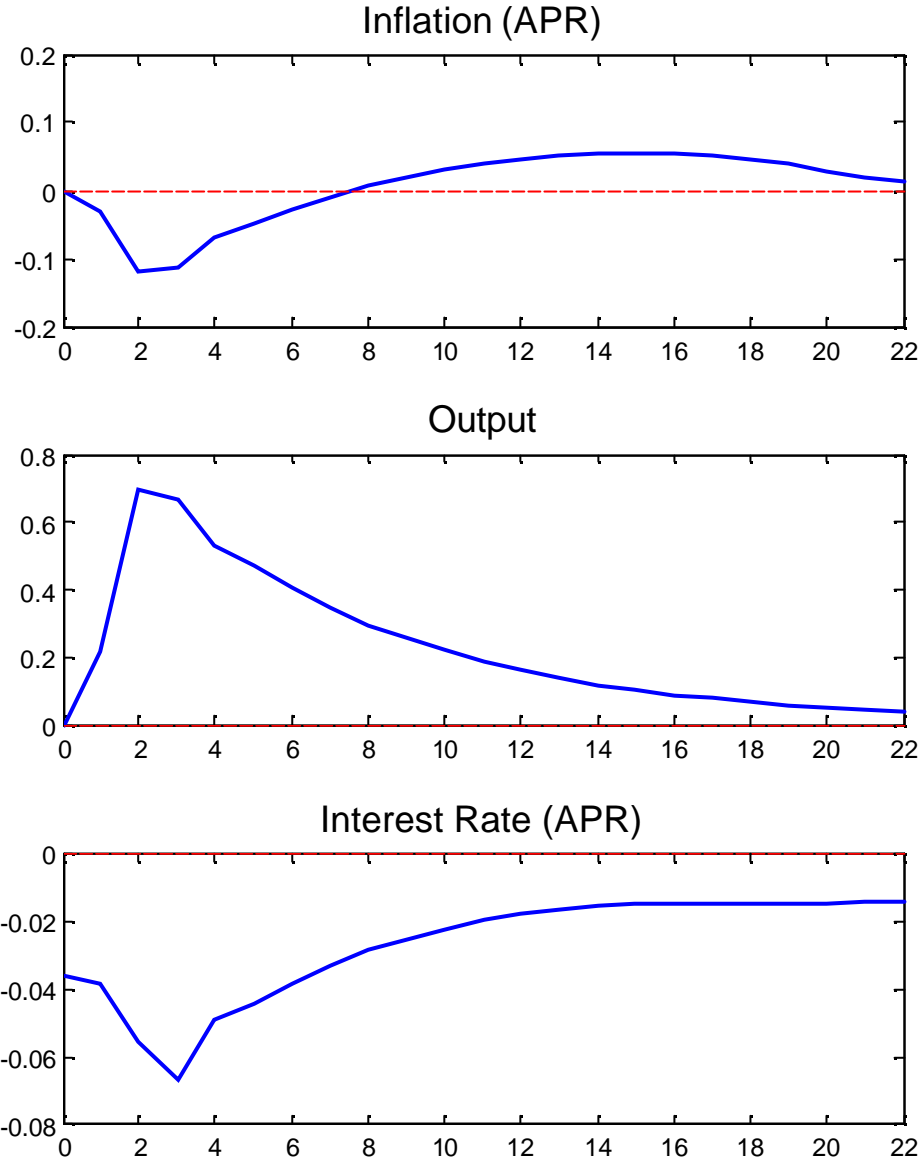


Figure 5: Interest Rate Shock (75 bps)

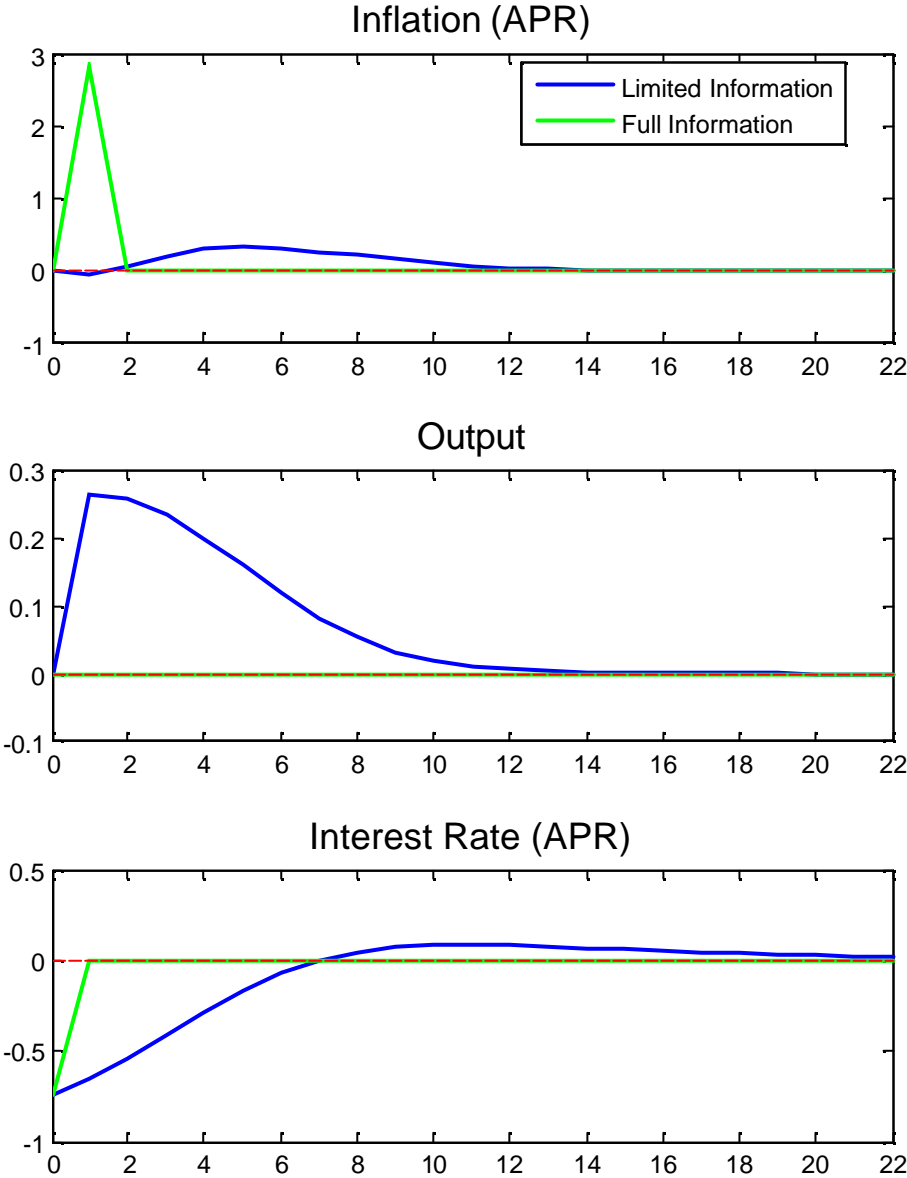


Figure 6: Inflation Target Shock (1 per cent)

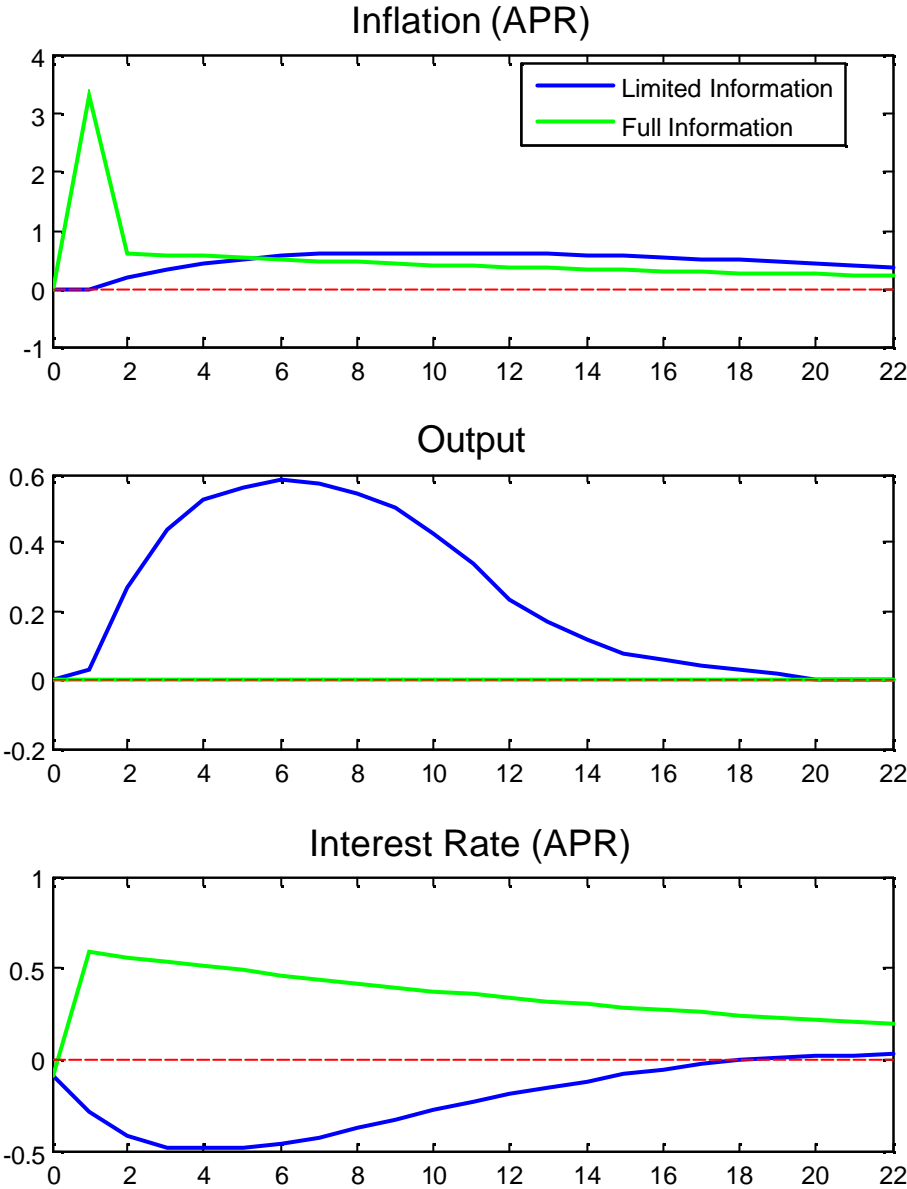


Figure 7: Central Bank Signal Shock (1 per cent)

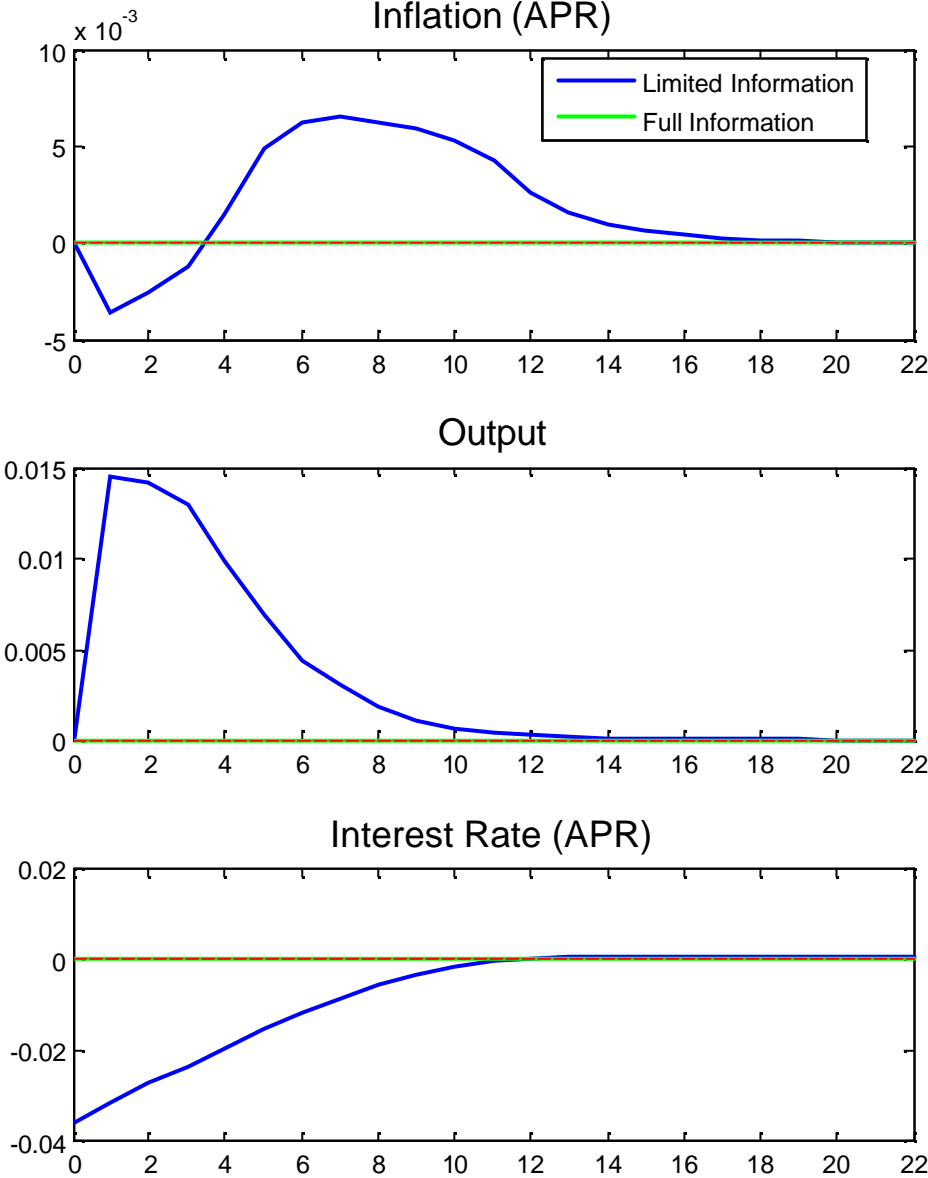


Figure 8: Technology Shock (1 per cent)

