Information, heterogeneity and market incompleteness in the stochastic growth model*

Liam Graham† and Stephen Wright‡

5 August 2007

Abstract

We provide a microfounded account of imperfect information in the stochastic growth model which dramatically changes the properties of the model. We describe heterogenous households that acquire information about aggregates through their participation in markets. If markets are incomplete, household information will be imperfect. We solve the model taking account of the infinite regress of expectations that this lack of information implies. We derive analytical and numerical results to show that imperfect information can significantly change the properties of the model: under virtually all calibrations the impact response of consumption to a positive aggregate technology shock is *negative*.

JEL classification: D52; D84; E32.

Keywords: imperfect information; higher order expectations; Kalman filter; dynamic general equilibrium

---

*We thank Kristoffer Nimark for invaluable comments and advice. We are also grateful for useful comments from Roland Meeks, Morten Ravn and seminar participants at the EUI Florence, the Institut für Weltwirtschaft, Kiel, the Indian Statistical Institute, Delhi, and the Universities of Bristol and Cambridge.

†Corresponding author: Department of Economics, University College London, Gower Street, London WC1E 6BT, UK. Liam.Graham@ucl.ac.uk

‡Department of Economics, Birkbeck College, University of London, Malet Street, London W1E 7HX, UK. s.wright@bbk.ac.uk

1 Introduction

Underlying most dynamic general equilibrium modelling is the assumption that households can perfectly observe the state variables. Complete markets rationalize this assumption: in a decentralized equilibrium households learn about aggregates through participating in markets, so if markets are complete so too will be information. However if markets are incomplete, households will in general be imperfectly informed about the aggregate economy, and hence about other agents. This means that rational households have to form expectations of aggregate states, and of other households' behaviour, leading to an infinite regress of expectations (Townsend, 1983, Woodford, 2002, Nimark 2007a,b).

We describe a version of the stochastic growth model in which households are heterogenous because they face an idiosyncratic productivity shock in addition to the standard aggregate productivity shock. If capital is the only tradeable asset, households' information is limited to a knowledge of their own capital holdings, along with the returns they observe from participating in labour and capital markets. We describe such households solving a signal extraction problem using a version of the Kalman filter that allows for endogeneity of the states (Baxter, Graham and Wright, 2007) and explicitly model higher-order expectations using techniques developed by Nimark (2007a).

A key feature of this paper is that the informational problem arises endogenously. There is no "noise" in our model, and the only assumption we make about the observability of aggregates is that households gain information about them through their participation in markets.

We derive analytical results which show that the economy with incomplete markets must differ from the full information economy, and show that the difference arises from a different response to aggregate productivity shocks. We further show that consumption in our economy is in general not certainty-equivalent (in the sense of Pearlman et al, 1986 or Svennson and Woodford, 2002, 2004). We then study the model numerically and find it differs dramatically from the complete markets economy.

In the standard stochastic growth model (which is a complete-markets version of our economy), the impact effect of a positive aggregate productivity shock on consumption is positive[1]. However, with incomplete markets and a consistent treatment of information we find:

1. Under a wide range of calibrations the impact of a positive aggregate productivity shock on aggregate consumption is negative.
2. The subsequent path of aggregate consumption is radically different from the full information case.
3. We show that certainty equivalence is a very good approximation over a wide range of calibrations.
4. This does not mean higher-order expectations are redundant since they improves the state forecasts which determine certainty-equivalent consumption.

The intuition for why our model behaves so differently under incomplete markets is as follows. Households only gain information about aggregates through the capital and labour markets in which they participate, so a household observes a positive innovation to aggregate productivity as positive innovations in its wage and the return to capital. The strong empirical evidence (for example Guvenen, 2005, 2007) that the variance of idiosyncratic productivity shocks is much higher than that of aggregate shocks means that the wage contains little useful information about aggregates, so the main signal the household receives is a positive innovation to the return to capital.

With imperfect information, households know that such a positive innovation to returns could be caused either by a positive innovation to aggregate productivity or by aggregate capital being lower than the household had previously estimated. The certainty equivalent response to the first would be to increase consumption, to the second to decrease consumption.

The relative weight of these two effects depends on the structure of the economy and the properties of the exogenous processes. But we show analytically that the second effect will, under reasonable parameter restrictions, always cause the consumption response to be less than under full informa-

[1] Campbell (1994) gives precise conditions under which the impact response of consumption is positive. A sufficient condition is that the coefficient of relative risk aversion is greater than unity.

tion; and we show numerically that under a wide range of calibrations the impact response of consumption to the shock is negative.

The study of imperfect information has a long history in macroeconomics. Here we pick out two strands of the recent literature which are particularly relevant to our work; a more complete review can be found in Hellwig (2006). The first strand looks at the problem of setting monetary policy under imperfect information. Most such models (Pearlman et al, 1986, Svennson and Woodford, 2002, 2004, Aoki, 2003, 2006) look at the problem of asymmetric information when the monetary policymaker has imperfect information but the private sector is perfectly informed. Pearlman (1992) and Svennson and Woodford (2003) look at the case where the private sector and the policymaker share the same imperfect information set.

A second strand, closer to the present paper, investigates the implications of the private sector having imperfect information. Keen (2004) investigates a model in which the private sector is poorly informed about the behaviour of the monetary policymaker and concludes that it can account for several business cycle features better than the standard model. In Collard and Dellas (2006), the introduction of imperfect information in an otherwise standard dynamic new Keynesian model can generate inflation persistence, hump-shaped dynamics of inflation and output and a liquidity effect. The effect of noise in productivity is investigated by Bomfim (2001) who shows that permanent / transitory confusion can lead to interesting business cycle dynamics and Lorenzoni (2006) who shows that measurement error in aggregate productivity can explain the presence of demand shocks.

In neither of these strands of the literature is an account given of the source of the informational restrictions: either measurement error is introduced ad hoc, or some variables are simply assumed not to be observed. In contrast, in our model the informational problem on the part of agents arises from the presence of an idiosyncratic productivity shock[2], and model information in a market-consistent matter, so information is only available to households through the markets in which they trade. We show that "noise" is not necessary to motivate informational problems.

Our households know that all other household in the economy face a

[2] Lorenzoni (2006) presents a model very close to ours in that households face both an aggregate and an idiosyncratic productivty shock. However the results of his paper depend on households observe a noisy signal of aggregate productivity, and this noise is simply assumed.

similar inference problem, but with a different information set. To forecast aggregates, a household must forecast the behaviour of all other households, which requires us to model higher-order expectations. Townsend (1983) first analysed the problem of "forecasting the forecasts of others" and the infinite regress of expectations that results. Woodford (2002) shows that the dynamics of such higher-order expectations can lead to shocks having more persistent effects. Here we draw on recent work by Nimark (2007a) who derives new techniques for modelling the resulting infinite-dimensional state vector when agents make dynamic choices.[3]

The literature on the relation between informational imperfections and incomplete markets is vast, but has mainly focussed on the implications for financial markets (Marin and Rahi, 2000, review some of this literature). In terms of the macroeconomy, Levine and Zame (2002) ask "Does market incompleteness matter?" and answer that it does not. Our paper shows that, while incomplete markets alone do not change aggregate properties, once we take account of the informational implications of incomplete markets the answer can change radically.

Section 2 describes the model and section 3 considers two benchmark cases. In section 4, we formalize the information set of agents, show how the infinite hierarchy of expectations arises and define the equilibrium. Section 5 presents our analytical results, and section 6 gives numerical results. We draw out some implications of our model in section 7 and conclude in section 8.

2 The stochastic growth model with incomplete markets and idiosyncratic productivity shocks

The economy consists of a large number of households and a large number of firms. We divide our economy into S islands, on each of which there are many firms and households. Households consume, rent capital and labour to firms and are subject to an island-specific shock to labour productivity. Firms use capital and labour to produce a single consumption good with a technology that is subject to an aggregate productivity shock. Markets

[3] Nimark's (2007b) analysis of the inference problem with higher expectations in a model of sticky prices is also a rare example of a paper in which, as in ours, imperfect information is due to heterogeneity rather than arbitrary noise.

are incomplete in that the only asset available is capital, and while labour is heterogenous across islands, we assume that capital is homogenous and can freely flow between islands. Our model is similar to versions of the stochastic growth model with heterogeneity by Maliar and Maliar (2003) and Nakajima (2005). We focus here on the key structural relationships; the full log-linearised model is provided in Appendix A

We use upper case letters for levels, lower case letters for log deviations. Letters without a time subscript indicate steady states. A superscript s indicates a variable relating to a household or firm on island s. Without the superscript the variable is an aggregate.

2.1 Households

A typical household on island s consumes (C_t^s) and rents capital (K_t^s) and labour (H_t^s) to firms. Household labour on each island has idiosyncratic productivity (Z^s) whereas capital is homogenous, so households earn the aggregate return (R_t) on capital but an idiosyncratic wage (V_t^s) on their labour. In our central case we assume that markets are incomplete and the only asset available to households is capital.

The problem of a household of type s is to choose paths for consumption, labour supply and investment (I^s) to maximize expected lifetime utility given by

$$E_t^s \sum_{i=0}^{\infty} \beta^i \left[\ln C_{t+i}^s + \theta \frac{\left(1 - H_{t+i}^s\right)^{1-\gamma}}{1-\gamma} \right] \tag{1}$$

where $\frac{1}{\gamma}$ is the intertemporal elasticity of labour supply, and β the subjective discount rate, subject to a resource constraint

$$R_t K_t^s + V_t^s H_t^s = C_t^s + I_t^s \tag{2}$$

and the evolution of the household's holdings of capital

$$K_{t+1}^s = (1-\delta) K_t^s + I_t^s \tag{3}$$

The expectations operator for an individual household is defined as the expectation given the household's information set Ω_t^s, i.e. for some variable q_t

$$E_t^s q_t = E_t q_t | \Omega_t^s$$

We assume that, apart from the idiosyncratic shock, households are identical across types and hence are unconditionally identical, and on any given island all households behave identically.

The household's first-order conditions consist of an Euler equation

$$C_t^s = \beta E_t^s \left[(1 + R_{t+1}) C_{t+1}^s \right] \tag{4}$$

and a labour supply relation

$$(1 - H_t^s)^{-\gamma} = \frac{V_t^s}{C_t^s} \tag{5}$$

2.2 Firms

The production function of a typical firm on island s is

$$Y_t^s = (J_t^s)^{1-\alpha} (A_t Z_t^s H_t^s)^a \tag{6}$$

where H_t^s is the labour of households on island s that the firm hires, J_t^s is capital and A_t is an aggregate productivity shock.

The first-order conditions of this firm are

$$R_t = \frac{Y_t^s}{J_t^s} \tag{7}$$

$$V_t^s = \frac{Y_t^s}{H_t^s} \tag{8}$$

where V_t^s is the wage paid to the labour that the firm hires.

2.3 Aggregates

Aggregate quantities are sums over household or firm quantities, and for convenience we calculate them as quantities per household type. For example aggregate consumption is given by

$$C_t = \frac{1}{S} \sum_{s=1}^{S} C_t^s \tag{9}$$

The economy's aggregate resource constraint is then

$$Y_t = C_t + I_t \tag{10}$$

Total capital equals the sum of both household and firm capital holdings:

$$K_t = \frac{1}{S}\sum_{s=1}^{S} K_t^s = \frac{1}{S}\sum_{s=1}^{S} J_t^s \tag{11}$$

but in general $J_t^s \neq K_t^s$, given free flow of capital between islands.

2.4 Markets

For simplicity, we assume that labour markets are completely segmented between islands, so firms on island s only rent labour from households on island s, and the wage on island s, V_t^s, adjusts to set labour supply (5) equal to labour demand (8).

In contrast, capital is assumed to be homogenous and tradeable between islands, so flows to islands with more productive labour. The aggregate return to capital, R_t, adjusts to clear the capital market, making the demand for capital for each firm (7) consistent with each household's Euler equation (4) and the aggregate resource constraint (10).

2.5 Shocks

For both the aggregate and idiosyncratic productivity shocks we assume autoregressive processes in deviations

$$a_t = \phi_a a_{t-1} + \varpi_t \tag{12}$$

$$z_t^s = \phi_z z_{t-1}^s + \varpi_t^s \tag{13}$$

where ϖ_t and ϖ_t^s are i.i.d mean-zero errors, and $E\varpi_t^2 = \sigma_a^2$; $E(\varpi_t^s)^2 = \sigma_z^2$. Following Campbell (1994), we assume that aggregate technology has a steady state growth rate of g.

We further assume that the innovation to the idiosyncratic process satisfies an adding up constraint, $\sum_{s=1}^{S} \varpi_t^s = 0$ which implies

$$\sum_{s=1}^{S} z_t^s = 0. \tag{14}$$

2.6 The system

While our underlying model is non-linear, we work with the log-linear approximation to the model which will allow us to use a linear filter to model the household's signal extraction problem[4].

We show in Appendix A.4 that the features of the economy relevant to a household of type s can be written as an Euler equation

$$E_t^s \Delta c_{t+1}^s = E_t^s r_{t+1} \tag{15}$$

and a linearised law of motion for the economy that is symmetric across types:

$$W_{t+1}^s = F_W W_t^s + F_c c_t + F_s c_t^s + v_t^s \tag{16}$$

where $W_t^s = \begin{bmatrix} \xi_t' & \chi_t^{s\prime} \end{bmatrix}'$ is a vector of underlying states relevant to a household of type s comprising aggregate states $\xi_t = \begin{bmatrix} k_t & a_t \end{bmatrix}'$ and states specific to household s, given by$\chi_t^s = \begin{bmatrix} \kappa^s & z_t^s \end{bmatrix}'$ where $\kappa_t^s = k_t^s - k_t$.The coefficient matrices F_W, F_c and F_s are defined in Appendix A.4.

The linearisation is very close to that Campbell (1994): indeed the coefficients for the aggregate part of our economy are identical to his.

Definition 1 ***(Equilibrium)*** *A competitive equilibrium for the above economy is a sequence of plans for*

- allocations of households $\left\{c_t^s, h_t^s, k_{t+1}^s\right\}_{t=1:\infty}^{s=1:S}$
- prices $\{r_t, w_t^s\}_{t=1:\infty}^{s=1:S}$
- aggregate factor inputs $\{k_t, h_t\}_{t=1:\infty}$

such that

1. Given prices and informational restrictions, the allocations solve the utility maximization problem for each consumer
2. $\{r_t, w_t^s\}_{t=1:\infty}^{s=1:S}$ are the marginal products of aggregate capital and labour of different types
3. All markets clear

[4]Log linearisation is common to all the literature on imperfect information cited in the introduction.

3 Benchmark cases

The main focus of this paper is an economy in which the only tradeable asset is capital, and, consistent with a decentralized equilibrium, agents are not directly provided with information on the aggregate states. However, as benchmark cases we first investigate the case of complete markets, which we show reveal full information, and that of incomplete markets with full information simply assumed.

Definition 2 ***(Full information)*** *Full information, which we denote by an information set Ω_t^*, is knowledge of the aggregate states in the economy ξ_t, the idiosyncratic states χ_t^s of all household types and the time-invariant parameters and structure of the underlying model Ξ.*

$$\Omega_t^* = \left[\xi_t, \{\chi_t^s\}_{s=1}^S, \Xi\right] \tag{17}$$

3.1 Complete markets

Complete markets imply the existence of a set of securities that span the distribution of idiosyncratic shocks, and that are freely available for all agents to trade. Thus complete risk-sharing is possible[5] and in the process the household productivity shocks z^s are revealed to all households, so each household knows both the aggregate wage and the full set of disaggregate wages. Risk-sharing implies that household paths of consumption are perfectly correlated so each household also knows aggregate consumption. Since households observe both the return to capital and the aggregate wage, it is straightforward to show that they can recover the aggregate state variables ξ_t.

Thus complete markets reveal complete information, and there is a representative household whose consumption is equal to aggregate consumption, which is a function only of the aggregate states:

$$c_t^* = \eta_\xi^{*\prime}\xi_t \tag{18}$$

where η_ξ' is a vector of time-invariant coefficients, see (C.13), that can be found by standard solution techniques for rational expectations models.[6]

[5] The net effect of the payoffs on these securities for each individual will be to replace the left-hand side of (2) with a constant share of aggregate income.

[6] Maliar and Maliar (2005) show that a complete markets economy with a closely related

Given perfect risk-sharing and symmetry all households have consumption equal to aggregate consumption.

3.2 Full information and incomplete markets

In this second special case we revert to our central assumption that the only asset available to agents is capital, so agents will be unable to trade away idiosyncratic risk. However we assume that despite the absence of markets that reveal the idiosyncratic states, agents nonetheless have full information, provided as an endowment. We show later (Proposition 3) that this assumption of incomplete markets and complete information is fundamentally inconsistent, but it nonetheless provides a useful analytical building block.

The properties of this economy are summarised in the following proposition:

Proposition 1 ***(Full information and incomplete markets)*** *In the heterogeneous economy with incomplete markets each household's optimal consumption under full information will take the form*

$$c_t^s|\Omega_t^* = \eta_W^{*\prime} W_t^s = \begin{bmatrix} \eta_\xi^{*\prime} & \eta_\chi^{*\prime} \end{bmatrix} \begin{bmatrix} \xi_t \\ \chi_t^s \end{bmatrix} \tag{19}$$

1. *The coefficients in η_ξ^* are identical to those under complete markets in equation (18).*

2. *The coefficients in η_χ^* solve the undetermined coefficients problem for η_ξ^* in a parallel complete markets economy in which the persistence of aggregate productivity is the same as that of idiosyncratic productivity ($\phi_a = \phi_z$) and the elasticity of intertemporal substitution is zero.*

3. *Aggregate consumption in this economy is identical to the complete markets solution in (18)*

$$\frac{1}{S}\sum_{s=1}^{S} c_t^s|\Omega_t^* = \eta^{*\prime}\xi_t = c_t^* \tag{20}$$

form of heterogeneity leads to a representative consumer with a utility function with "preference shocks". This does not arise in our economy due to the adding up constraint across idiosyncratic shocks (14), and the multiplicative nature of the shocks (a case noted by Maliar and Maliar (2005) in their footnote 2).

Proof. *See Appendix C* ■

Corollary 1 *Under full information, the idiosyncratic element in consumption is a random walk, and the idiosyncratic element in capital is a unit root process*

The combination of incomplete markets and complete information (provided as an endowment) results in an economy which is identical at an aggregate level to the complete markets economy, but which differs markedly at a household level.

A household's response to an idiosyncratic productivity shock is very different from its response to an aggregate productivity shock because idiosyncratic productivity shocks do not affect expectations of future returns, so are, from the household's point of view, simply a change in permanent income. Proposition 1 states that these optimising responses to idiosyncratic shocks will be identical to what the responses of aggregate consumption to aggregate shocks would be in a complete markets economy with no intertemporal substitution. The permanent income response to idiosyncratic shocks in turn implies that the idiosyncratic component of consumption is a random walk as in Hall (1978).[7] However, the adding-constraint across idiosyncratic shocks means that such permanent shifts in idiosyncratic consumption cancel out in the aggregate.

4 Incomplete markets and imperfect information

Our basic assumption is that markets are incomplete in the sense that capital is the only tradeable asset. Our key assumption in all that follows that households only obtain information from the markets they participate in so, consistent with our assumption on markets, we can write the information set of a household of type s at time t as

$$\Omega_t^s = \left[\{r_i\}_{i=0}^t , \{w_i^s\}_{i=0}^t , \{k_i^s\}_{i=0}^t , \Xi \right] \tag{21}$$

[7] Recall Campbell's (1994) result that such an economy will generate consumption responses in line with the permanent income hypothesis.

where Ξ contains the parameters and structure of the underlying model and is therefore time-invariant[8].

We define a measurement vector $i_t^s = \begin{bmatrix} r_t & w_t^s & k_t^s \end{bmatrix}'$ such that the information set evolves according to

$$\Omega_{t+1}^s = \Omega_t^s \cup i_{t+1}^s \tag{23}$$

and we can write the measurement vector as

$$i_t^s = H_w W_t^s + H_c c_t \tag{24}$$

where the matrices H_w and H_c are defined in Appendix A.5.

We show in the Appendix A.5 that the three observables are given by

$$\begin{aligned} r_t &= \lambda_3 (a_t + h_t - k_t) && (25) \\ k_t^s &= \kappa_t^s + k_t && (26) \\ w_t^s &= w_t + z_t^s && (27) \end{aligned}$$

so this information set does not, in general, allow households to recover either aggregate or idiosyncratic states.

The informational problem in our model arises because, since aggregates are not directly observable, households are unable to distinguish between aggregate and idiosyncratic productivity shocks. Thus innovations in the observable variables could be caused either by true innovations to the exogenous processes, or by households' estimates of the aggregate states being incorrect.

4.1 The hierarchy of expectations

In this section we define the state vector relevant to household s, consisting of the underlying states W_t^s (defined after (16)) and an infinite hierarchy of average expectations (Townsend, 1983, Woodford, 2002, Nimark, 2007a) of

[8] Households also have knowledge of the history of their own optimising decisions, defined as

$$\left[\{c_i^s\}_{i=0}^{t}, \{h_i^s\}_{i=0}^{t-1} \right] \tag{22}$$

however, since each of these histories embodies the household's own responses to the evolution of Ω_t^s, it contains no information not already in Ω_t^s.

the same vector. The consumption of a household of type s is then given by

$$c_t^s = \eta' E_t^s X_t^s \tag{28}$$

where the (infinite dimension) state vector is

$$X_t^s = \left[\begin{array}{ccccc} W_t^s & W_t^{(1)} & W_t^{(2)} & W_t^{(3)} & .... \end{array} \right]' \tag{29}$$

The first-order average expectation $W_t^{(1)}$ is an average over all households' expectations of their idiosyncratic state vector

$$W_t^{(1)} = \frac{1}{S}\sum_{s=1}^{S} E_t^s W_t^s \tag{30}$$

and higher-order expectations are given by

$$W_t^{(k)} = \frac{1}{S}\sum_{s=1}^{S} E_t^s W_t^{(k-1)}; k > 1 \tag{31}$$

Given the household consumption function (28), the definition of aggregate consumption (9) implies

$$c_t = \eta' X_t^{(1)} \tag{32}$$

where

$$X_t^{(1)} = \left[\begin{array}{cccc} W_t^{(1)} & W_t^{(2)} & W_t^{(3)} & .... \end{array} \right]' \tag{33}$$

In the case of full information, described in Section 3, the hierarchy becomes redundant, since higher-order elements in the hierarchy are simply equal to the true values of the averages across non-expectational states, given by

$$W_t^{(k)} | \Omega_t^* = \left[\begin{array}{c} \xi_t \\ 0 \end{array} \right] : k \geq 1 \tag{34}$$

where the lower block of zeros exploits the adding up constraint in (14).

4.1.1 A heuristic argument

To see why X_t^s is the relevant state vector, consider the following heuristic argument. First assume that households think that only the non-expectational states W_t^s affect their own consumption, so that

$$c_t^s = \eta_1' E_t^s W_t^s$$

Assuming the household knows that all other households will behave in the same way (we formalize this in Assumption 1 below), aggregate consumption will be

$$c_t = \frac{1}{S}\sum_{s=1}^{S} c_t^s = \eta_1' W_t^{(1)}$$

But then the original consumption function is mis-specified since to correctly forecast the aggregate economy using (16) household s must forecast aggregate consumption, which depends on $W_t^{(1)}$.[9] So a better specification is

$$c_t^s = \eta_2' E_t^s \begin{bmatrix} W_t^s \\ W_t^{(1)} \end{bmatrix}$$

However this implies that aggregate consumption will be

$$c_t = \eta_2' \begin{bmatrix} W_t^{(1)} \\ W_t^{(2)} \end{bmatrix}$$

hence the household state vector should again be augmented to include $W_t^{(2)}$, and so forth. This leads to an "infinite regress" (Townsend, 1983) of expectations, so the state vector of individual s must contain an infinite hierarchy of expectations.

4.2 The household's signal extraction problem

To implement optimal consumption (28), a household of type s must form estimates of the infinite-dimension state vector X_t^s by using the information Ω_t^s available to it. The optimal linear filter is the Kalman filter, however this problem differs significantly from the standard Kalman filter in two ways. The first difference is that the states depend on the household's choice variable c_t^s. Baxter, Graham and Wright (2007) describe this "endogenous" Kalman filter in detail, and gives conditions for its stability and convergence which are satisfied here. Secondly, since the aggregate states depend on

[9] We show below that the trivial case where $W_t^{(1)} = \frac{1}{S}W_t^s = \begin{bmatrix} \xi_t' & 0 \end{bmatrix}'$ cannot hold, nor in general can $W_t^{(1)}$ be expressed as a simple linear combination of non-expectational states, hence it is indeed required as an additional state vector

aggregate consumption, and hence the behaviour of all other households, we need to make an assumption about what household s knows about the behaviour of all other households. We follow Nimark (2007) in assuming that each household applies the Kalman Filter to the entire model on the assumption that each other household is behaving in the same way.

Assumption 1: *It is common knowledge that all households' expectations are rational (model consistent).*

Nimark (2007) discusses this assumption in more detail, but it is essentially a generalization of the full information rational expectations assumption.

Given Assumption 1, we show in Appendix D that each household faces a symmetric signal extraction problem of the form

$$X^s_{t+1} = Lc^s_t + MX^s_t + Nv^s_{t+1} \quad (35)$$

$$i^s_t = H'X^s_t \quad (36)$$

where L, M, N and H are matrices yet to be determined and i^s_t is the measurement vector of household s, defined before equation (23).

Proposition 2 ***(Equilibrium with market-consistent imperfect information)*** *In an economy in which each household*

a. Has an information set of the form (21)

b. Forms optimal forecasts of the states X^s_t by solving the household-specific filtering problem given by (35) and (36)

c. Chooses consumption to satisfy its Euler equation (15)

An equilibrium which satisfies Assumption 1 and Definition 1 is a fixed point of the following undetermined coefficients problem:

$$M = \left\{ \begin{bmatrix} F_W & F_c\eta' \\ 0_{\infty,r} & L\eta' + (I - \beta H')M \end{bmatrix} + \begin{bmatrix} 0_{r,\infty} \\ \beta H'M \end{bmatrix} \right\} \quad (37)$$

$$N = \begin{bmatrix} I_r \\ \beta HN \end{bmatrix} \quad (38)$$

$$\eta' = (\eta' - R')\left[M + L\eta'\right] \quad (39)$$

where β, the gain matrix of the endogenous Kalman filter, is defined in Appendix D and shows how the measured variables update the state estimates

$$E_t^s X_t^s - E_{t-1}^s X_{t-1}^s = \beta \left(i_t^s - E_t^s i_t^s \right) \tag{40}$$

and L, H and R are defined in Appendix D.

Proof. See Appendix D ■

4.3 Solution technique

We solve the iterative system of equations given by (37) to (39) for a typical household. The solution to this problem implies a law of motion both for any individual household's state estimates, which evolve by (40), but also, when we average across such updating rules, for the hierarchy of average expectations. This in turn, via (32), determines the solution for aggregate consumption, consistent with each household solving a symmetric filtering and optimal consumption problem. While we model the behaviour of a typical household, there is no representative household in this economy.

5 Properties of the economy with incomplete markets and imperfect information

In this section we derive analytical results which show that imperfect information changes the nature of the economy, and explain the mechanism behind this. We further show that consumption in our model is in general not certainty equivalent, but that we can decompose household consumption into a certainty-equivalent response and a component arising from the hierarchy of expectations. To simplify the analysis, the propositions in this section consider only the case of fixed labour supply ($\gamma = \infty$).[10]

Proposition 3 ***(Non-Replication of Full Information) If the variance of the idiosyncratic shocks is non-zero*** ($\sigma_z > 0$)***, the economy described in Proposition 2 can never replicate the full information***

[10] We conjecture that all remain valid under variable labour supply. This can be verified numerically, but the analytical proofs become much more convoluted. In the proofs we note the implications of relaxing the assumption.

economy. However deviations from full information are transitory even when there are permanent shocks to underlying states, and the informational problem does not change the steady state.

Proof. See Appendix E ■

In the economy we describe, households have a restricted information set, given by (21), which arises from the factor markets they trade in. Proposition 3 shows that this informational problem always matters for the equilibrium of the economy. The proof makes clear that this result is non-trivial, and that the idiosyncratic productivity shocks cause the informational problem. We show in our numerical results, Section 6, that the differences from the full information equilibrium are quantitatively significant, but Proposition 3 states that such differences must be transitory. Imperfect information can only have implication for business cycle dynamics, and not for the long-run growth properties of the model.

Corollary 2 ***As the economy approaches the limiting homogeneous case (as*** $\sigma_z \to 0$***) it approaches the complete markets economy. Furthermore, in this limiting case, as*** $t \to \infty$ ***the entire history of returns*** $\{r_s\}_{s=1}^{t}$ ***becomes informationally redundant.***

Proposition 3 draws out the crucial link between household heterogeneity, incomplete markets and information. In the limit, with no idiosyncratic shocks, all households are, and know themselves to be, identical. Even though markets are incomplete, this only matters to the extent that households differ from each other. Since risk-sharing markets can only smooth out the impact of idiosyncratic shocks, their absence becomes unimportant as these idiosyncratic shocks disappear, as do the informational problems associated with incomplete markets. In the limit each household can perfectly observe both the aggregate wage and the return, and thereby trivially infer the values of the aggregate states. But the corollary goes further than this: given a sufficiently large number of observations, households do not even need the history of returns: an information set consisting only of the history of aggregate wages is sufficient to reveal the states.

5.1 The impact of aggregate productivity shocks

We have shown that the economy with imperfect information must differ from the full information economy. Since the adding-up constraint across idiosyncratic shocks (14) means that the aggregate economy is only driven by the process for aggregate productivity, the differences from full information must arise from a different dynamic response to aggregate productivity shocks. The following proposition states the key features of this response.

Proposition 4 ***(Impact effects of aggregate productivity shocks)*** *In the economy characterized by Proposition 3, a positive aggregate productivity shock has the following effects on impact:*

a) Household estimates of aggregate capital unambiguously fall;

b) If the persistence of idiosyncratic productivity is less than some value strictly greater than the persistence of aggregate productivity (i.e. $\phi_z < \bar{\phi}_z > \phi_a$) household (and hence aggregate) consumption is unambiguously lower than under full information.

Proof. See Appendix F. ■

In our economy households must base their estimates of underlying states on the signals they observe from markets. When a positive aggregate productivity shock hits, each household will observe this as a simultaneous rise in the aggregate return on capital and their own wage. While the former is a "pure" signal of aggregates, the latter also contains information on idiosyncratic states. As such it can be interpreted as a “noisy” signal of the aggregate economy, although it differs from standard signal-noise problems in that here what is noise with respect to the aggregate economy conveys information about idiosyncratic states that is also important to the household.

The first part of this proposition states that the signal extraction problem means that a positive aggregate productivity shock always causes households to revise *downwards* their estimate of aggregate capital. So what is unambiguously good news under full information becomes, under incomplete markets, a mixture of good and bad news, causing the consumption response to be smaller on impact.

To see why the estimate of capital falls, note that a general property of optimal filtering is that forecasts of states must always have lower variance than the actual states[11]. With respect to aggregate productivity this implies that the household's estimate must respond less to shocks than does actual

[11] For some variable q_t and household s's estimate thereof, $E_t^s q_t$ we can write

$$q_t = E_t^s q_t + f_t^s$$

productivity. For a productivity shock in period 1, and assuming we start from the steady state, this means $E_1^s a_1 < a_1$. Estimates must be consistent with observations, i.e. $E_1^s r_1 = r_1$. The return to capital is given from (25) by $r_t = \lambda_3 (a_t - k_t)$ so

$$a_1 - k_1 = E_1^s (a_1 - k_1) \tag{41}$$

Since capital is predetermined, $k_1 = 0$ so

$$E_1^s k_1 = E_1^s a_1 - a_1 < 0 \tag{42}$$

Thus the estimate of capital must fall on the impact of a positive innovation to aggregate productivity.

The nature of the consumption response is also driven by the requirement that state estimates are consistent with observations. Households know their own capital, which is predetermined. This implies that if a household revises its estimate of aggregate capital downwards, it must revise its estimate of the idiosyncratic component of its own capital ($\kappa_t^s = k_t^s - k_t$) *upwards* by exactly the same amount. It is also quite easy to show (see Appendix G) that the same must apply for the estimate of idiosyncratic productivity. Thus bad news on capital in the aggregate economy is always offset by good news on the idiosyncratic economy.

As idiosyncratic productivity becomes more persistent, an estimated positive innovation to idiosyncratic productivity becomes better news. But the parameter restriction in part b) of Proposition 4 states that, unless the persistence of idiosyncratic productivity becomes very high, the bad news about aggregate capital will always outweigh the good news on the idiosyncratic economy. Since aggregate shocks affect all households symmetrically (though not observably so), this implies that the response of aggregate consumption must also be strictly less than under full information.

We show numerically in Section 6.6 below that $\bar{\phi}_z$, the upper bound for ϕ_z, the persistence of idiosyncratic productivity, is always very close to unity, so this is very close to being a general result.

where f_t^s is a filtering error. Efficiency of the filter implies $cov\left(q_t, f_t^s\right) = 0$ so

$$var\left(q_t\right) \geq var\left(E_t^s q_t\right)$$

5.2 A reduced state vector and implications for consumption.

The equilibrium described by Proposition 3 requires each household to form an optimal forecast of all elements of the hierarchy of average expectations. However, the following proposition states that the dimension of this problem can be significantly reduced:

Proposition 5 ***(Reduced state vector)*** *A sufficient state vector for the aggregate economy is* $\left[\begin{array}{cccc} \xi_t' & a_t^{(1)} & a_t^{(2)} & ... \end{array} \right]'$

Proof. See Appendix G ■

This result arises because, for each household, estimates of the non-expectational states, $E_t^s W_t^s$ must, to be consistent with the information set, satisfy the following adding up constraint

$$H_W' \left[W_t^s - E_t^s W_t^s \right] = 0 \tag{43}$$

which arises directly from the measurement equation (24). We show in Appendix G that this implies restrictions across the full hierarchy of average expectations and leads to the reduced state vector. This in turn implies an alternative specification of the consumption function in (28) which will prove useful in subsequent analysis

Corollary 3 ***(Decomposition of consumption function)*** *The consumption function for household s can be written in the form*

$$c_t^s = \eta_W^{*\prime} E_t^s W_t^s + \sum_{k=1}^{\infty} \mu_k E_t^s \left[a_t - a_t^{(k)} \right] \tag{44}$$

5.3 Certainty equivalence

It is a standard result[12] in the existing literature on optimising behaviour under symmetric imperfect information that the property of certainty equivalence holds: optimal choices are the same linear function of estimated state variables as of actual state variables under full information. In the context of our model this implies the following definition:

[12] See, for example, Pearlman et al (1986), Svensson and Woodford (2002, 2004); Baxter Graham and Wright (2007)

Definition 3 ***(Certainty Equivalence):*** *Each household's consumption function is certainty-equivalent if it can be expressed in the form*

$$c_t^s = \eta_W^{*\prime} E_t^s W_t^s$$

where η_W^* is the vector of coefficients in the consumption function under full information in Proposition 1.

We showed in Section 4.1 that this property will not hold in our model, and optimal consumption will instead depend on the full hierarchy of expectations. But by inspection of (44), this dependence can be broken down into a certainty-equivalent response, and a response which is driven by the extent to which the household believes that the hierarchy of expectations of aggregate technology differs from the household's own estimate. Using this decomposition we can show that the concept of certainty equivalence still plays an important role in our model.

Proposition 6 ***(Deviations from Certainty Equivalence)***

1. *The two limiting cases of the economy, as σ_s tends to zero (the homogeneous case), and as $\sigma_s \to \infty$ (extreme heterogeneity), are both certainty equivalent.*

2. *For intermediate cases certainty equivalence does not hold.*

Proof. See Appendix H ∎

Corollary 2 means that the limiting homogeneous case is trivially certainty equivalent. To see why the limiting case of extreme heterogeneity is also certainty equivalent, we need to consider the link between market incompleteness and informational problems.

While returns provide a signal exclusively about the aggregate block of the economy, for the general case the household's wage w_t^s provides a signal about both aggregate and idiosyncratic blocks. However, as agents become more heterogeneous the signal from the wage is increasingly dominated by the impact of the idiosyncratic economy. As σ_s tends to infinity, the economy is effectively segmented into two distinct blocks, with returns providing the only information about the aggregate block, and the wage providing information only about the idiosyncratic block. Each household updates estimates of aggregate states using only information on returns, which is common knowledge, so from Assumption 1 each household knows that every other household will update their estimates in the same way. Hence all households have identical estimates of aggregate states, which straightforwardly implies that the entire hierarchy of expectations of aggregate states is known, and equal to each household's estimates. Hence, by inspection of (44), certainty equivalence must hold in this limiting case.

6 Numerical results

6.1 Calibration

Given the degree of uncertainty over some of the key parameters, the values we choose in this section should be seen as giving a baseline case which we use to generate impulse response functions and give intuition for our results. We carry out sensitivity analysis to all of the important parameters in Section 6.6 below.

The key parameters are the persistence and innovation variance of the aggregate and idiosyncratic productivity processes. We calibrate the aggregate productivity shock with the benchmark RBC values for persistence of $\phi_a = 0.9$ and an innovation standard deviation $\sigma_a = 0.7\%$ per quarter (Prescott, 1986). In Appendix B we discuss the details of our calibration of the idiosyncratic technology process, drawing on the empirical literature on labour income processes. It turns out that a calibration which sets idiosyncratic persistence equal to aggregate persistence (i.e. $\phi_z = \phi_a = 0.9$) is consistent with Guvenen's (2005, 2007) recent estimates using US panel data. There is however strong evidence that idiosyncratic technology has a much higher innovation standard deviation. In Appendix B we show that a figure of 4.9% per quarter is consistent with Guvenen's results.

Card (1994) estimates the intertemporal elasticity of labour supply, $\frac{1}{\gamma}$ to be between 0.05 and 0.5. For our baseline calibration, we choose $\gamma = 5$, in the middle of this range. For the other parameters we follow Campbell (1994)[13].

[13] $\delta = 0.025; \alpha = 0.667; \beta = 0.99; H = 0.33; g = 0.005$

6.2 Numerical solution method

All of our theoretical results relate to a representation with an infinite dimension state vector. Nimark (2007a) shows that the infinite hierarchy can be approximated to an arbitrary accuracy by a finite order representation. We adapt his approach by truncating the hierarchy and writing a state vector of the form

$$\bar{X}_t^s = \left[\begin{array}{cccc} W_t^s & W_t^{(1)} & .... & W_t^{(h)} \end{array} \right]' \tag{45}$$

where h is the order of the truncation. For our baseline calibration, we use $h = 5$. Adding an extra order to the hierarchy would change the impact effect of consumption reported below by 10^{-7}.

6.3 The nature of impulse response functions

Before discussing impulse responses in our calibration, we should note an important caveat. The response profiles discussed in the next three sections differ from standard impulse response functions under full information, in that we examine the impact of shocks to the two underlying stochastic processes, a_t and z_t^s that are unobservable to any agents in the economy. The impulse response functions we obtain could not therefore be observed contemporaneously.

As a result of this informational asymmetry between agents and the observer, the stochastic properties of the model are crucial in determining the nature of impulse response functions, in a way that they are not under full information. Under full information, after the initial shock has taken place, the remainder of the impulse response is equivalent to a perfect foresight path, and is thus known in advance to both observer and agents in the model. In contrast, under incomplete information, the agents in the model are continuously making inferences from new information as it emerges, and thus can only imperfectly predict their future behaviour. In making these inferences the underlying stochastic properties of the model are crucial, in a way that they are not under full information.[14]

[14] To be precise, impulse responses under incomplete information depend on the parameters in the true covariance matrix of structural shocks Q, whereas under full information they do not.

6.4 Response to an aggregate productivity shock

Figure 1 shows the response of consumption in our baseline model to a 1% positive innovation in the process for aggregate productivity. This clearly demonstrates our key result: the response of aggregate consumption is significantly negative on impact of a positive productivity shock.

Figure 1: Response of consumption to a 1% positive innovation to aggregate productivity[15]

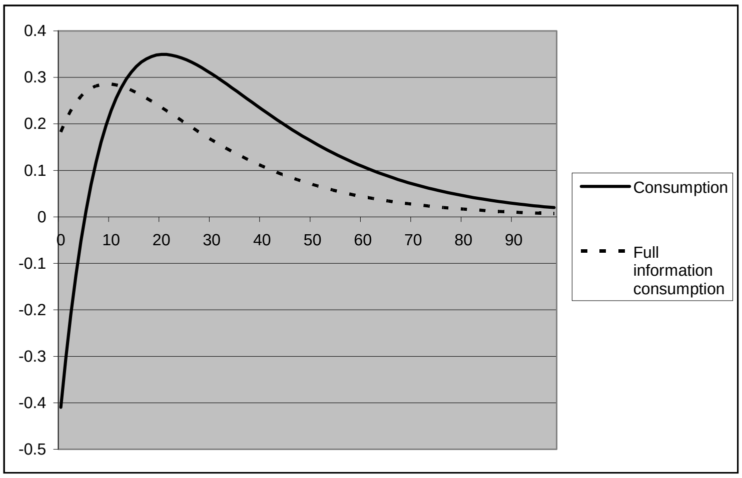

For comparison, Figure 1 also shows the response of consumption under full information which mimics the standard RBC result. Under full information consumption increases on impact by 0.18%, under imperfect information it falls by 0.41%. In Section 6.6 below we show that this negative response is robust to a very wide range of variation in structural parameter values. Here we give an intuitive explanation of this feature.

On impact, a household does not observe the aggregate productivity shock directly, but only the associated positive innovations to the market return and the idiosyncratic wage. The household knows that these innovations could have arisen either because there were structural shocks this period, or because the state estimates on which its previous forecasts were based were incorrect. In response to the innovations, the household updates its state estimates using (40) and it is the revised state estimates that determine the response of consumption response via (28).

[15] x-axis shows periods; y-axis shows percentage deviations from steady state

Corollary 3 shows that the response of the economy can be decomposed into two components: a certainty-equivalent response, and a response dependent on the hierarchy of expectations. Idiosyncratic productivity in our baseline calibration is much more volatile than aggregate productivity, and numerically we can show that the impulse responses are close to those in the limiting case of extreme heterogeneity. So the responses are dominated by the certainty-equivalent part, and the impact of the hierarchy is quantitatively small. We discuss each of the components of the response in turn.

6.4.1 The certainty-equivalent response

Assume for purposes of illustration that the economy is initially in steady state in period 0, before the shock occurs, with all states equal to their steady state values of zero. Given this initial position, under the assumption of certainty equivalence, the idiosyncratic consumption function (44) reduces to

$$c_1^s = \begin{bmatrix} \eta_k & \eta_a & \eta_\kappa & \eta_z \end{bmatrix} \begin{bmatrix} E_1^s k_1 \\ E_1^s a_1 \\ E_1^s \kappa_1^s \\ E_1^s z_1^s \end{bmatrix} \tag{46}$$

Since all agents are identical aggregate consumption is equal to idiosyncratic consumption $c_t = c_t^s \; \forall t$,(but crucially, given the caveats of Section 6.3, not observably so).

Figure 2: Response of state estimates to a 1% positive innovation to aggregate productivity

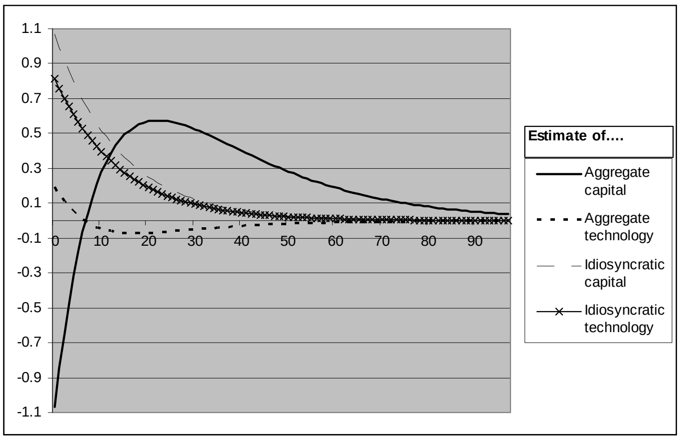

Figure 2 shows how these state estimates respond to the innovations in the return and the idiosyncratic wage caused by the aggregate productivity shock. To understand these responses, consider the impact on each of the state estimates in turn

1. Aggregate productivity (a_t). We have shown in Proposition 4 that the estimated value must increase by strictly less than the true increase. Numerically we can show that even the sign of the response is ambiguous. A positive innovation in returns could be caused by a positive innovation to aggregate productivity this period, but it could also be due to capital having been previously over-estimated. Since agents already know what the return was in previous periods, the only way this given historic return profile can be reconciled with a lower estimate of capital both now and in the past is if they also revise down their estimate of past levels of productivity . The more persistent is aggregate productivity, the greater this offsetting effect will be. In our baseline calibration the two effects are of similar magnitude so that, as shown in Figure 2, the impact of the innovations on the household's estimate of aggregate productivity is very close to zero.

2. Aggregate capital (k_t). Proposition 4 shows that the estimate of aggregate capital must fall on impact. Given the observed positive innovation to returns, the less a given increase in true aggregate technology is attributed to estimated technology, the worse the news must be for aggregate capital: i.e. the more it is interpreted by households as a signal that their estimates of aggregate capital were too high in the past.

3. The Idiosyncratic component of capital (κ_t^s). As we noted in our discussion of Proposition 4, since households observe their own capital directly, any change in their estimate of aggregate capital must be precisely offset by an updated estimate of the idiosyncratic element of their own capital (i.e. $\kappa_t^s = k_t^s - E_t^s k_t$).

4. Idiosyncratic productivity (z_t^s). An increase in the wage always causes households to increase their estimates of idiosyncratic productivity. The more heterogeneous the economy the more a rise in the wage is attributed to this cause.

In the Appendix C we show $\eta_k > \eta_\kappa$ and for the parameter restriction in Proposition 4 (i.e. $\phi_z < \overline{\phi}_z > \phi_a$) $\eta_a > \eta_z$. In our calibration, $\eta_k = 0.57$, $\eta_\kappa = 0.08; \eta_a = 0.18$ and $\eta_z = 0.08$. Since $E^s_t k_1 < 0$ and $E^s_t a_1 \approx 0$, this means that, given (46) and the changes in the state estimates described above, consumption falls on impact.

In the next period the household again observes innovations to its idiosyncratic wage and the market return and these will again differ from the forecasts since the household's estimates of the states were different from the true states. The household repeats the process described above, of updating its state estimates using the information contained in the innovations and then using these new estimates to form forecasts of the observed variables. In the next period, the realisations of these variables will again differ from the forecasts, giving the household more information about the true states. Given the initially low level of consumption compared to full information, the actual capital stock is higher throughout, and hence as state estimates improve consumption ultimately overshoots the full information response before ultimately returning to the steady state.

6.4.2 The impact of the hierarchy

The calibrated case turns out to be numerically very close to the limiting case of extreme heterogeneity. This means that households are very close to having common estimates of aggregate states, and hence, via (44) the consumption function is very close to being certainty-equivalent. Table 1 shows the impact effects on the different orders of the hierarchy.

Table 1: Impact effect of an aggregate technology shock on the hierarchy of expectations

i	1	2	3	4	5
$a_1^{(i)}$	0.1896	0.1737	0.1734	0.1734	0.1734

Given Proposition 6, which states that at the two limiting cases consumption will be certainty equivalent, there will be some value of the variance of the idiosyncratic shock which maximises the deviation from certainty equivalence, and hence the impact of the hierarchy. Numerically we find that this value is very small, but since this in turn implies that deviations from full information are small the quantitative impact of the hierarchy is always small. For variances one tenth of the value in our baseline calibration, the deviation from certainty equivalence remains quantitatively unimportant.

The limited nature of deviations from certainty equivalence does not imply that the hierarchy of expectations is redundant. Even when certainty equivalence is close to holding in terms of state estimates, these estimates themselves are more efficient due to the improved forecasts each household can make of the economy.

6.5 Response to an idiosyncratic productivity shock

With incomplete markets, households cannot insure against idiosyncratic shocks. Under full information, the household knows that an idiosyncratic shock has no effect on aggregates and hence none on future returns, so the response follows the permanent income hypothesis, as in Proposition 1.

Figure 3: Response of idiosyncratic consumption to a 1% positive innovation to idiosyncratic productivity

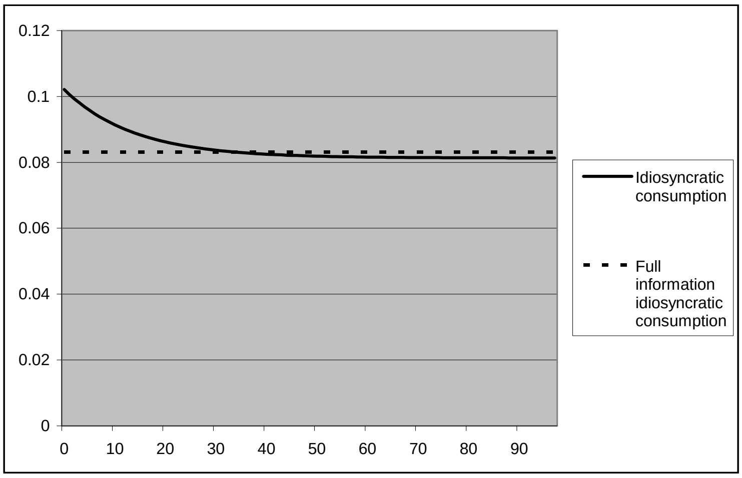

Figure 3 shows the response a given household to a 1% positive innovation to the process for idiosyncratic productivity. With imperfect information, consumption overshoots the full information response on impact, and then converges only slowly back. The response is very close to that under full information, in contrast to the case of an aggregate productivity shock where imperfect information dramatically changes the response. This is a consequence of the large variance of the idiosyncratic shock compared to the aggregate shock: the observed innovations are unconditionally much more likely to be due to an idiosyncratic shock. [16]

Although the response to idiosyncratic shocks is small (0.08% compared to a full information impact response of 0.26% for an aggregate productivity shock of the same magnitude), their much higher variance means that the behaviour of each household is dominated by these shocks and the macroeconomy is in effect a sideshow.

6.6 Sensitivities

Our key result is that, under imperfect information, the response of aggregate consumption to a positive aggregate productivity shock is always less than under full information, and it is negative in our baseline calibration. In this section we examine how robust this is to changes in the calibration.

Apart from the elasticity of labour supply, $\frac{1}{\gamma}$,the standard real business cycle parameters do not have any great effect on the result, since while they change the structure of the economy they do not change the nature of the informational problem that drives our result. Since the informational problem is about identifying whether an aggregate or idiosyncratic shock has occurred, the properties of these two processes impact our results.

Proposition 4, part b) states there is a threshold value $\bar{\phi}_z$ of the persistence of the idiosyncratic shock for which the impact response is less than that under full information. Table 2 shows this threshold, both for the fixed labour supply case considered in the proposition and the calibrated value of γ, for different values of the persistence of aggregate technology. Clearly, Proposition 4b is very close to being a general result: consumption under responds when compared with the full information case.

[16] As the economy tends to the limiting case of extreme heterogeneity we can show that household responses to idiosyncratic shocks will precisely match those under full information - only the response to aggregate shocks will be distorted. Conversely, as the economy approaches homogeneity ($\sigma_z \to 0$), and thus, from Corollary 2, gets arbitrarily close to replicating the full information response to an aggregate shock, the household's over-response to an idiosyncratic shock attains a maximum. But this is a rational response, given the vanishingly small contribution of idiosyncratic shocks to innovations in the household's information set.

Table 2: Critical values, $\bar{\phi}_z$ (as defined in Proposition 4) of persistence of idiosyncratic shock

ϕ_a	0.95	0.9	0.8	0.5	0.2	0
Fixed labour: $\gamma = \infty$	0.998	0.997	0.995	0.994	0.993	0.993
Variable labour:$\gamma = 5$	0.997	**0.996**	0.995	0.993	0.993	0.992

NB: base case shown in bold

Table 3 shows how the impact response of consumption to a true aggregate productivity shock varies with the persistence of the aggregate shock, ϕ_a and that of the idiosyncratic shock, ϕ_z (the baseline calibration is in bold). Unconditional variances determine the signal extraction problem, so as the persistences fall, so too does the degree of the informational problem and the response of consumption becomes less negative. However for the idiosyncratic process, there is an offsetting effect As the idiosyncratic shock becomes more persistent, the "good news" from an estimated innovation to idiosyncratic productivity offsets the "bad news" on aggregate capital, so the response of consumption becomes less negative. As ϕ_z approaches the critical values in table 2 the response of consumption becomes less negative.

Table 3: Impact effect of aggregate technology shock on aggregate consumption: sensitivity to persistence parameters

	ϕ_a				
ϕ_z	0.95	0.9	0.85	0.7	0.5
0.95	−0.541	−0.338	−0.238	−0.115	−0.059
0.9	−0.614	**−0.410**	−0.301	−0.157	−0.087
0.85	−0.603	−0.423	−0.318	−0.172	−0.098
0.7	−0.426	−0.374	−0.305	−0.180	−0.107
0.5	−0.097	−0.245	−0.241	−0.169	−0.107

NB: base case shown in bold

Table 4 shows how the impact response of consumption to a true aggregate productivity shock varies as the innovation standard deviation σ_z and persistence of the idiosyncratic process (ϕ_z) change. The second column, with an infinite variance of the idiosyncratic shock, corresponds to the limiting heterogeneous case of Proposition 6, the final column, with a zero variance, to the limiting homogenous case.

Table 4: Impact effect of aggregate technology shock on aggregate consumption: sensitivity to properties of idiosyncratic shock

	σ_z/σ_a					
ϕ_z	∞	10	5	2	1	0
0.95	−0.352	−0.345	−0.338	−0.273	−0.113	0.183
0.9	−0.440	−0.425	**−0.410**	−0.276	0.022	0.183
0.85	−0.474	−0.448	−0.424	−0.211	0.058	0.183
0.7	−0.510	−0.438	−0.376	−0.009	0.126	0.183
0.5	−0.526	−0.365	−0.245	0.0763	0.160	0.183

NB: base case shown in bold

As the relative standard deviation of the idiosyncratic shock decreases, going from left to right in the table, the information problem becomes less acute so the impact response of consumption becomes less negative. As the persistence of the shock falls, the unconditional variance falls so the informational problem becomes less acute. However this is offset by the second effect described above. Since the unconditional variance is a multiple of the innovation variance, the relative strength of the first effect depends on the magnitude of the innovation variance. For high values of the innovation variance the second effect is dominant. For values in the middle of the variance range, the first effect dominates for low values of persistence, and the second effect for high values.

Tables 3 and 4 show that our result is robust to any reasonable calibration of the productivity shocks. Only for relatively non-persistent idiosyncratic shocks with standard deviations around five times lower than those estimated in the literature, does the impact response of consumption come close to that under full information.

Table 5: Impact effect of aggregate technology shock on aggregate consumption: sensitivity to the elasticity of labour supply,

γ	∞	20	10	5	2	1	0.5
	−0.272	−0.306	−0.341	**−0.410**	−0.601	−0.936	−1.55

NB: base case shown in bold

Finally we examine sensitivity of our results to the elasticity of labour supply. The left-most column, with $\gamma = \infty$, corresponds to the case of fixed labour supply. As labour supply becomes more elastic, moving right, consumption responds more negatively. To understand why this is so, note that returns depend on aggregate labour supply $r_t = \lambda_3 (a_t + h_t - k_t)$. When labour supply becomes more elastic, the household varies it more in response to the observed innovations to returns and the wage. Since all households behave identically in response to an aggregate productivity shock, this means aggregate labour increases by more (though not observably so to any individual household), and returns increase by more on impact, so the observed innovation to returns becomes greater. The negative response of consumption arises from the ambiguity in the signal provided by a positive innovation to returns, so, as labour supply responds more, consumption will respond more negatively.

7 Conclusions

We believe that our model is only a starting-point for the analysis of the link between heterogeneity, market incompleteness and informational problems. We have shown a very stark contrast with the standard complete markets model; but we do not yet know how robust this contrast will be to further modifications.

On the one hand it might easily be argued that capital is the only asset is too drastic a deviation from the standard model, given that we do observe at least some risk-sharing by financial markets. Introducing a limited, if still incomplete range of tradeable financial assets could change our results.

Also, markets are not the only source of information available to households: government statistical offices provide estimates of aggregates which must in some degree ease the informational problem. If we were to model a statistical office that could a) measure the true factor incomes of individual households; b) sample a large number of household types; and c) could do so in real time, its estimates of total output would approach the true value, and it would be straightforward in our log-linear economy to infer the true values of both aggregate states, allowing each household in turn to back out values of idiosyncratic states. But no actual statistical office would claim to be able to resolve all practical problems associated with a) and b); and the practical implementation of c) is essentially impossible. Furthermore the problem of inferring states from output would be enormously harder in a complex, non-linear economy. Hence, while a statistical office could certainly mitigate the informational problem it cannot solve it.[17]

On the other hand, while our results show that the informational limitations due to incomplete markets can have very significant effects it is very easy to argue that we may be significantly *understating* the extent of the informational problem. Our model is highly simplified, with only a single source of idiosyncratic uncertainty and strong assumptions of symmetry across households. A central assumption of our analysis is that agents do at least know the structure and parameters of our model; and our application of the Kalman Filter makes the common assumption that the filtering parameters have converged. There is a large body of research, both on model uncertainty (for example Hansen and Sargent, 2001) and learning (Evans and Honkapohja, 2003) that would question these assumptions. In the context of our model, a natural question to ask is whether the joint time series process for the observables that arises from the solution to the filtering problem is learnable; and if it is, whether the underlying structural parameters of the model are uniquely identified. Even if both these strong conditions are met, it is easy to see that the inferences made by agents in our model would require very large amounts of data.

Until these issues have been investigated, we would hesitate to draw strong empirical conclusions from our analysis. Nonetheless our results do show a very distinct contrast from the standard benchmark model in breaking, or at least weakening, the positive short-term correlation between consumption, employment, and underlying returns on capital, implied by full information. We suspect that the alternative dynamics implied by our analysis may in due course generate insights into the well-known puzzles in macroeconomics and finance relating to these correlations.

[17] Hayek's (1945) original analysis of the signalling role of prices was embedded in a critique of the distorting role of state planning. Our analysis suggests that ,even in a fully competitive economy there may after all be some role for the state, albeit of a limited nature, in mitigating the informational problems caused by incomplete markets.

References

Abowd, J.M. and David Card, (1989), "On the Covariance Structure of Earnings and Hours Changes," Econometrica 57, pp 411-445

Aiyagari, S Rao, (1994) "Uninsured Idiosyncratic Risk and Aggregate Saving," Quarterly Journal of Economics 109, 3, pp 659-84,

Aoki, Kosuke, (2003), "On the optimal monetary policy response to noisy indicators", Journal of Monetary Economics 50 (3), pp 501-523.

Aoki, Kosuke, (2006), "Optimal commitment policy under noisy information", Journal of Economic Dynamics and Control 30, pp 81-109.

Baxter, Brad, Liam Graham and Stephen Wright (2007), "The Endogenous Kalman Filter", Working Paper

Bomfim, Antulio N, (2001), "Measurement error in general equilibrium: the aggregate effects of noisy economic indicators", Journal of Monetary Economics 48, pp 505-603.

Campbell, John Y, (1994), "Inspecting the mechanism: an analytical approach to the stochastic growth model", Journal of Monetary Economics 33, pp 463-506.

Card, David, (1994), "Intertemporal Labor Supply: An Assessment", in Sims, Christopher (1994), "Advances in Econometrics, Sixth World Congress".

Chang, Yongsung and Sun-Bin Kim, (2006), "From Individual To Aggregate Labor Supply: A Quantitative Analysis Based On A Heterogeneous Agent Macroeconomy", International Economic Review 47, 1, pp1-27

Collard, Fabrice and Harris Dellas, (2006), "Imperfect Information and Inflation Dynamics", Working Paper.

Evans, George and Seppo Honkapohja, (2003), "Learning and Expectations in Macroeconomics", Princeton University Press

Guvenen, Fatih, (2005) "An Empirical Investigation of Labor Income Processes", Working Paper

Guvenen, Fatih, (2007), "Learning Your Earning: Are Labor Income Shocks Really Very Persistent?", American Economic Review 97 (3), pp 687-712.

Hall, Robert E, (1978), "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence", Journal of Political Economy 86(6), pp 971-87

Hamilton, James, (1994), "Time Series Analysis", Princeton University Press

Hansen, Lars and Tom Sargent, (2001), "Robust Control and Model Uncertainty", Working paper

Hayek, FA, (1945), "The use of knowledge in society", American Economic Review 35 (4), pp. 519-530

Hellwig, Christian, (2006), "Monetary Business Cycle Models:Imperfect Information", UCLA Working Paper

Keen, Benjamin D, (2004), "The Signal Extraction Problem Revisited:Its Impact on a Model of Monetary Policy", Working Paper.

Levine, David and William Zame, (2002), "Does market incompleteness matter?", Econometrica 70 (5), pp 1805-1839

Lorenzoni, Guido, (2006), "A Theory of Demand Shocks", NBER Working Paper 12477

MaCurdy, T, (1982): The use of time series processes to model the error structure of earnings in a longitudinal data analysis, Journal of Econometrics,18, pp 83-114.

Maliar, Lilia and Serguei Maliar, (2003), "The representative consumer in the neoclassical growth model with idiosyncratic shocks, Review of Economic Dynamics, 6 pp 362–380

Marin, Jose and Rohit Rahi, (2000), "Information Revelation and Market Incompleteness", Review of Economic Studies, 67(3), pp. 563-579

Nakajima, Tomoyuki, (2005),"A business cycle model with variable capacity utilization and demand disturbances",European Economic Review 49 pp 1331-1360

Nimark, Kristoffer, (2007a), "Dynamic Higher Order Expectations", Working Paper

Nimark, Kristoffer, (2007b), "Dynamic Pricing and Imperfect Common Knowledge", Journal of Monetary Economics, forthcoming.

Pearlman, Joseph, (1992), "Reputational and nonreputational policies under partial information", Journal of Economic Dynamics and Control 16 (2), pp 339-357.

Pearlman, Joseph, David Currie, and Paul Levine,1986, ”Rational Expectations Models with Partial Information”, Economic Modeling 3, pp 90-105.

Prescott, Edward C, (1986), "Theory ahead of business cycle measurement", Federal Reserve Bank of Minneapolis Quarterly Review 10 (4), pp 9-22.

Svennson, Lars and Michael Woodford, (2002), "Optimal policy with partial information in a forward-looking model: Certainty-equivalence redux", Princeton University Working Paper.

Svennson, Lars and Michael Woodford, (2004), "Indicator variables for optimal policy under asymmetric information.", Journal of Economic Dynamics and Control 28 (4), pp.661-690.

Townsend, Robert M, (1983), "Forecasting the Forecasts of Others", Journal of Political Economy 91 (4), pp 546-588.

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 Honour of Edmund S. Phelps, Princeton:Princeton University Press.