ELSEVIER

Volume 57, issue 2
March 2010
ISSN 0304-3932

Journal of MONETARY ECONOMICS

CONTENTS

Published by ELSEVIER
in collaboration with the
WILLIAM E. SIMON GRADUATE SCHOOL OF BUSINESS ADMINISTRATION, UNIVERSITY OF ROCHESTER, and the DEPARTMENT OF ECONOMICS, BOSTON UNIVERSITY

Journal of Monetary Economics 57 (2010) 164–174

Contents lists available at ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jme

Information, heterogeneity and market incompleteness ☆

Liam Graham [a,*], Stephen Wright [b]

[a] Department of Economics, University College London, Gower Street, London WC1E 6BT, UK
[b] Department of Economics, Birkbeck College, University of London, Malet Street, London W1E 7HX, UK

A R T I C L E I N F O

Article history:
Received 31 October 2008
Received in revised form
23 December 2009
Accepted 23 December 2009
Available online 4 January 2010

JEL classification:
D52
D84
E32

Keywords:
Imperfect information
Higher order expectations
Kalman filter
Dynamic general equilibrium

A B S T R A C T

Information is "market-consistent" if agents only use market prices to infer the underlying states of the economy. This paper applies this concept to a stochastic growth model with incomplete markets and heterogeneous agents. The economy with market-consistent information can never replicate the full information equilibrium, and there are substantial differences in impulse responses to aggregate productivity shocks. These results are robust to the introduction of a noisy public signal and aggregate financial markets. We argue that the principle of market-consistent information should be applied to any model with incomplete markets.

1. Introduction

A number of recent papers (notably Krusell and Smith, 1998) argue that market incompleteness does not much matter for the aggregate properties of DSGE models. This paper shows that, once account is taken of its informational consequences, market incompleteness *does* matter.

Most DSGE models assume complete information, i.e. knowledge of the current state vector. In the stochastic growth model this consists of aggregate capital and aggregate technology, quantities which are intrinsically difficult to measure directly. Radner (1979) shows that complete markets resolve this measurement problem. If markets are complete the information set consisting of market prices is invertible, i.e. it reveals the states. The assumptions on markets and information are consistent with each other.

In contrast, models with incomplete markets typically simply *assume* the states are directly observable. We propose that the concept of "market-consistent information" should be applied to models with incomplete markets. This paper analyses a linearized stochastic growth model with heterogeneous agents in which households have heterogeneous information sets arising from the limited set of markets they trade in. *Average* expectations then have persistent effects via aggregate capital, leading to a "hierarchy of expectations" (Townsend, 1983; Woodford, 2003; Nimark, 2007). The paper extends this methodology to allow for endogenous states and market-consistent information. This extension is a

☆ We are grateful to Robert King (the Editor) and an anonymous referee for comments that substantially improved the paper. We also thank Kristoffer Nimark for his advice, and numerous colleagues and participants in seminars for their input.

* Corresponding author.
E-mail addresses: Liam.Graham@ucl.ac.uk (L. Graham), s.wright@bbk.ac.uk (S. Wright).

0304-3932/$ - see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.jmoneco.2009.12.005

non-trivial exercise and our novel solution techniques allow the discipline of market-consistent information sets to be applied to any linearized DSGE model with incomplete markets.

This paper shows that a market-consistent rational expectations equilibrium will have the following three properties. First, the equilibrium will never replicate the full information solution. As a direct consequence, the good news of a positive productivity shock must always be interpreted at least in part as the bad news that the capital stock has been over-estimated in the past. Second, for virtually all empirically plausible calibrations this bad news response dominates the short and medium term response to a productivity shock. Third, this marked contrast with the complete markets response is robust to the inclusion of a noisy public signal and aggregate financial markets.

There is a long tradition of the study of imperfect information in macroeconomics, a review of which can be found in Hellwig (2006). The approach taken in this paper is distinct from recent work in explicitly drawing the link between information and markets; showing how the hierarchy of expectations can be modelled when there are endogenous states; and finding that imperfect information has significant effects in the stochastic growth model, without the need for noisy indicators as in Lorenzoni (2010) or strategic complementarities as in Angeletos and La'O (2009).

The remainder of this paper is structured as follows. Section 2 presents the model. Analytical results are in Section 3, and numerical results in Section 4. Section 5 discusses the robustness of these results, and Section 6 concludes.

2. The model

This section presents a model of the type that is becoming standard in the dynamic general equilibrium literature.[1] There are a large number of households and a large number of firms, divided into S islands, on each of which there many firms and households. There are shocks to aggregate and island-specific labour productivity.

Upper case letters denote levels, lower case letters denote log deviations from the steady state growth path. A superscript s indicates a variable relating to a typical 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_t^s) whereas capital is homogenous, so households earn the aggregate return (R_{kt}) on capital but an idiosyncratic wage (V_t^s) on their labour. Apart from the idiosyncratic shock, households on different islands are unconditionally identical.

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

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

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

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

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 a_t

$$E_t^s a_t = E a_t | \Omega_t^s \tag{4}$$

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

$$\frac{1}{C_t^s} = \beta E_t^s \left[\frac{R_{t+1}}{C_{t+1}^s} \right] \tag{5}$$

where $R_t = R_{kt} + 1 - \delta$ is the gross return to a one-period investment in capital, and a labour supply relation

$$\theta (1-H_t^s)^{-\gamma} = \frac{V_t^s}{C_t^s} \tag{6}$$

[1] Examples of papers which use similar models include Krusell and Smith (1998) and Lorenzoni (2010).

2.2. Firms

The production function of a typical firm on island s is

$$Y_t^s = (J_t^s)^{1-\alpha}(A_t Z_t^s H_t^s)^a \tag{7}$$

where A_t is an aggregate productivity shock and J_t^s is the capital rented by the firm: in general, $J_t^s \neq K_t^s$, since capital will flow to more productive islands. The first-order conditions of this firm are

$$R_{kt} = \frac{Y_t^s}{J_t^s}, \quad V_t^s = \frac{Y_t^s}{H_t^s} \tag{8}$$

2.3. Aggregates

Aggregate quantities are sums over household or firm quantities, and for convenience are calculated as quantities per household. For example aggregate consumption is given by $C_t = (1/S)\sum_{s=1}^{S} C_t^s$. The economy's aggregate resource constraint is then

$$Y_t = C_t + I_t \tag{9}$$

2.4. Markets

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 (6) equal to labour demand (8). In contrast, capital is 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 (8) consistent with each household's Euler equation (5) and the aggregate resource constraint (9).

2.5. Shocks

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

$$a_t = \phi_a a_{t-1} + \varpi_t \tag{10}$$

$$z_t^s = \phi_z z_{t-1}^s + \varpi_t^s \tag{11}$$

where ϖ_t and ϖ_t^s are i.i.d mean-zero errors, and $E\varpi_t^2 = \sigma_a^2$; $E(\varpi_t^s)^2 = \sigma_z^2$. The innovation to the idiosyncratic process satisfies an adding up constraint, $\sum_{s=1}^{S} \varpi_t^s = 0$ which implies

$$\sum_{s=1}^{S} z_t^s = 0 \tag{12}$$

2.6. Linearization

The remainder of the paper works with the log-linear approximation to the model in order to be able to use a linear filter to model the household's signal extraction problem.[2] Under a loglinear approximation, the features of the economy relevant to a household on island s can be written as an Euler equation

$$E_t^s \Delta c_{t+1}^s = E_t^s r_{t+1} \tag{13}$$

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

$$W_{t+1}^s = F_W W_t^s + F_c c_t + F_s c_t^s + u_t^s \tag{14}$$

where $W_t^s = [\xi_t' \quad \chi_t^{s'}]'$ is a vector of underlying states relevant to a household on island s comprising aggregate states $\xi_t = [k_t \quad a_t]'$ and states specific to the household, given by $\chi_t^s = [k_t^s - k_t \quad z_t^s]'$. The innovation vector for W_t^s is $u_t^s = [0 \quad \varpi_t \quad 0 \quad \varpi_t^s]'$ (both k_t and k_t^s are pre-determined). The structural coefficient matrices F_W, F_c and F_s are functions of the deep parameters of the model. The precise relations and the details of the linearization are provided in appendix B1, which is available as supplemental material on Science Direct.

[2] Although non-linear filters exist, incorporating one in this type of model and dealing with the hierarchy of expectations in a non-linear world would present a formidable challenge.

2.7. The economy with full information

If complete markets exist, risk-sharing implies that consumption is perfectly correlated across households so each household also knows aggregate consumption and the idiosyncratic component of its labour income. It is then easy to show that the information set is "instantaneously invertible"—full information can be recovered using only t-dated information.

Definition 1 (*Full information*). Full information, denoted 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 households and the time-invariant parameters and structure of the underlying model Ξ:

$$\Omega_t^* = [\xi_t, \{\chi_t^s\}_{s=1}^S, \Xi] \tag{15}$$

On the other hand, if markets are incomplete but full information is simply assumed, it is straightforward to show that the economy is identical at an aggregate level to the complete markets economy, but differs markedly at a household level.[3] This is related to Krusell and Smith's (1998) result that an economy with incomplete markets can closely resemble one with complete markets (the resemblance is exact in the model presented here because it is linearized). However, as what follows will show, this result is in general dependent on a market-*inconsistent* assumption of complete information.

3. Market-consistent information

With only capital and labour markets the market-consistent information set of a household on island s at time t is

$$\Omega_t^s = [\{r_i\}_{i=0}^t, \{\nu_i^s\}_{i=0}^t, \Xi] \tag{16}$$

where Ξ contains the parameters and structure of the underlying model.[4] Define a measurement vector $i_t^s = [r_t \;\; \nu_t^s]'$ such that the information set evolves according to $\Omega_{t+1}^s = \Omega_t^s \cup i_{t+1}^s$. Using the structural equations of the economy,

$$i_t^s = H_W' W_t^s + H_c c_t \tag{17}$$

where the matrices H_W and H_c are defined in Appendix B1.

Since aggregates are not directly observable, households are unable to distinguish between aggregate and idiosyncratic productivity shocks. They must therefore make errors (and will know that they must make errors) in estimating the true values of the states. 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. But the informational problem for each household is not restricted to forming estimates of the states, W_t^s, since each household knows that the aggregate capital stock depends on the average expectation of these states. This in turn must imply that the average expectation of the average expectation also matters, and so on—hence the true problem has an infinite dimensional state vector.

3.1. The hierarchy of expectations

Since in general the market-consistent information set will differ across households, the state vector relevant to household on island s, X_t^s, can be shown to consist of the non-expectational states W_t^s, defined after (14), and an infinite hierarchy of average expectations of W_t^s (Townsend, 1983; Woodford, 2003; Nimark, 2007, 2008)[5]

$$X_t^s = [W_t^s \;\; W_t^{(1)} \;\; W_t^{(2)} \;\; W_t^{(3)} \;\; \ldots]' \tag{18}$$

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

$$W_t^{(1)} = \frac{1}{S}\sum_{s=1}^{S} E_t^s W_t^s \tag{19}$$

and higher-order expectations are given by

$$W_t^{(k)} = \frac{1}{S}\sum_{s=1}^{S} E_t^s W_t^{(k-1)}, \quad k > 1 \tag{20}$$

[3] The permanent income response to idiosyncratic shocks implies that the idiosyncratic component of consumption is a random walk as in Hall (1978). However, the adding-up constraint across idiosyncratic shocks (12) means that such permanent shifts in idiosyncratic consumption cancel out in the aggregate. In general, this form of uninsurable income uncertainty would be expected to cause precautionary saving which would change the steady state of the model, but these effects are precluded by linearizing.

[4] Households also have knowledge of the history of their own optimizing decisions, $\{c_i^s\}_{i=0}^t$, $\{h_i^s\}_{i=0}^{t-1}$, $\{k_i^s\}_{i=0}^t$, 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. Note that while k_t^s is directly observable, the idiosyncratic component of capital, $k_t^s - k_t$, contained in the vector of underlying states, W_t^s, is not.

[5] The online Appendix (B1) provides a heuristic argument that demonstrates how the hierarchy of expectations arises from the problem of infinite regress, and illustrates why the Law of Iterated Expectations cannot be used to simplify this problem.

3.2. The household's filtering problem

To implement optimal consumption a typical household on island s must form estimates of the 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 from the standard Kalman filter in three ways. The first difference is that there is no extraneous noise. The second is that the states depend on the household's choice variable c_t^s. Baxter et al. (2009) describe this endogenous Kalman filter in detail, and give conditions for its stability and convergence which are satisfied here. Third, since the aggregate states depend on aggregate consumption, and hence the behaviour of all other households, an assumption is needed about what a household on one island knows about the behaviour of households on all other islands. 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).

This is essentially a generalization of the full information rational expectations assumption given idiosyncratic information sets.

Given Assumption 1, each household faces a symmetric endogenous Kalman filter problem of the form

$$X_{t+1}^s = Lc_t^s + MX_t^s + Nu_{t+1}^s \tag{21}$$

$$i_t^s = H'X_t^s \tag{22}$$

where L, M, N and H are matrices yet to be determined, u_t^s is the structural state innovation in (14) and i_t^s is the measurement vector of household s, defined in (17).

3.3. Equilibrium with market-consistent information

Definition 2 (*Rational expectations equilibrium with market-consistent information*). A competitive equilibrium in which the law of motion of the economy is consistent with each agent solving a decentralized optimizing and informational problem, i.e. a sequence of plans for allocations of households $\{c_t^s, h_t^s, k_{t+1}^s\}_{t=1:\infty}^{s=1:S}$; prices $\{r_t, v_t^s\}_{t=1:\infty}^{s=1:S}$; and 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. $\{r_t, v_t^s\}_{t=1:\infty}^{s=1:S}$ are the marginal products of aggregate capital and island-specific labour.
3. All markets clear.

Appendix A1 shows how an equilibrium satisfying Assumption 1 and Definition 2 can be derived. Although in general this problem can only be solved numerically, the following analytical result describes a key characteristic of the equilibrium.

Proposition 1 (*Non-replication of full information*). *In a non-homogenous economy* ($\sigma_Z > 0$) *with the market-consistent information set* (15) *an equilibrium satisfying Assumption* 1 *and Definition* 2 *never replicates the full information equilibrium.*

Proof. See Appendix A2.

This result contradicts the simple intuition that arises from counting shocks and observables. Each household needs to identify two underlying shocks, ϖ_t and ϖ_t^s, and has two observables, their (idiosyncratic) wage and the (aggregate) return on capital. Since the latter is only affected by the aggregate shock, ϖ_t, it might appear that ϖ_t could be recovered simply from innovations to returns. The proof shows that there is indeed a fixed point of the filtering problem in which this happens; but that this fixed point is unstable. In time series terms there is a reduced form ARMA representation of r_t in which ϖ_t is the innovation, but this representation is "non-fundamental" (Lippi and Reichlin, 1994), and hence ϖ_t cannot be recovered from the history of r_t.

The stability condition that is violated here is mathematically identical to that given by Fernandez-Villaverde et al. (2007) which must be satisfied for an econometrician to be able to infer true structural shocks and impulse responses from an estimated vector autoregressive representation of the economy in which the number of underlying shocks equals the number of observable variables. If this condition is not satisfied, Fernandez-Villaverde et al. show that the true shocks and impulse responses will be those to a "non-fundamental" representation of the economy. Given that this condition is not satisfied in the model economy, the typical household is faced with exactly the same inference problem as Fernandez-Villaverde et al.'s econometrician.

A corollary to this result helps understand the mechanism which makes the responses differ from those under full information. This corollary can be proved for the case of fixed labour supply ($\gamma = \infty$) and also in the neighbourhood of this case, i.e. for $\gamma > \overline{\gamma}$, where $\overline{\gamma}$ is finite. Numerically, the corollary appears always to hold.[6]

[6] Appendix A3 provides a sufficient condition for the average estimate of capital to be negative, which appears numerically always to be satisfied.

Corollary 1 (*Impact effects of aggregate productivity shocks*). *A positive aggregate technology shock leads households to reduce their estimate of aggregate capital*

Proof. See Appendix A3.

This result arises from the general property of optimal filters that state estimates have lower variance than states. The only source of aggregate fluctuations in the model is the technology shock. The solution to the filtering problem implies that aggregate technology must have lower variance than actual technology, so when such a shock hits, the estimate of technology must rise by strictly less than actual technology, i.e. starting from a steady state in $t=0$, $E_1^s a_1 < a_1$. Assuming fixed labour supply to simplify things, the linearized return is, as in Campbell (1994) $r_t = \lambda(a_t - k_t)$ and since estimates must be consistent with the information set, $E_1^s r_1 = r_1 = \lambda(a_1 - k_1)$, so $E_1^s k_1 - k_1 = E_1^s a_1 - a_1$. Since capital is predetermined, $k_1 = 0$ so $E_1^s k_1 = E_1^s a_1 - a_1 < 0$. Thus each household's estimate of capital (and hence the average estimate) must *fall* on the impact of a positive innovation to aggregate productivity. What is unambiguously good news under full information appears to the average household to be a mixture of good and bad news.

4. The response to aggregate productivity shocks

While the previous section shows qualitative differences from the full information equilibrium, this section investigates their quantitative implications.

4.1. Calibration

The key parameters are the persistence and innovation variance of the aggregate and idiosyncratic productivity processes.[7] Aggregate productivity is calibrated 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). The calibration of idiosyncratic technology draws on the empirical literature on labour income processes. A calibration that sets idiosyncratic persistence equal to aggregate persistence (i.e. $\phi_z = \phi_a = 0.9$) appears 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. A figure of 4.9% per quarter is consistent with Guvenen's results.[8]

4.2. The nature of impulse response functions

The response profiles discussed in this section differ from standard impulse response functions under full information in that they show the impact of a shock to an underlying stochastic process, a_t, that would be unobservable to any agents in the economy. These impulse response functions could not therefore be observed contemporaneously and the stochastic properties of the model are crucial in determining the nature of impulse response functions.

In contrast, 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. The agents in the model are continuously making inferences from new information as it emerges, and thus are uncertain not only about the value of future shocks, but also about their own future behaviour in response to past shocks.[9]

4.3. Response to an aggregate productivity shock

Fig. 1 shows the effect of a 1% positive innovation in the process for aggregate productivity on aggregate consumption in the baseline model and in the case of full information.[10] Under full information consumption increases on impact. In contrast, with incomplete markets and market-consistent information, the response of aggregate consumption is significantly negative on impact of a positive productivity shock.

This response is driven by the nature of the filtering problem. Households do not observe the aggregate technology shock directly, but only the associated positive innovations to the aggregate return and the idiosyncratic wage. They then use these observed innovations to update their estimates of the states, and it is these state estimates which determine their consumption decision.

[7] The values for the standard parameters $\delta = 0.025$; $\alpha = 0.667$; $\beta = 0.99$; $g = 0.005$ are chosen following Campbell (1994). The choice of log utility has already restricted the coefficient of relative risk aversion to unity. Card (1994) estimates the intertemporal elasticity of labour supply, $1/\gamma$ to be between 0.05 and 0.5. The calibration sets $\gamma = 5$, in the middle of this range. Steady state labour is taken to be $H = 0.33$ which implies the weight of labour in the utility function is $\theta = 3.5$.

[8] See Appendix B3 for a full discussion. Our calibration technique takes account of households' observing their own labour supply, but it can be argued that some idiosyncratic innovations to labour productivity may also be directly observable. This issue is discussed in Section 5.1.

[9] It is not the form of impulse responses that are unobservable, but the shocks that feed into them. The assumption of common knowledge of rationality means that any household could draw Fig. 1, but no household would be able to identify contemporaneously that a productivity shock had actually occurred.

[10] Appendix B4 gives details of the numerical solution method.

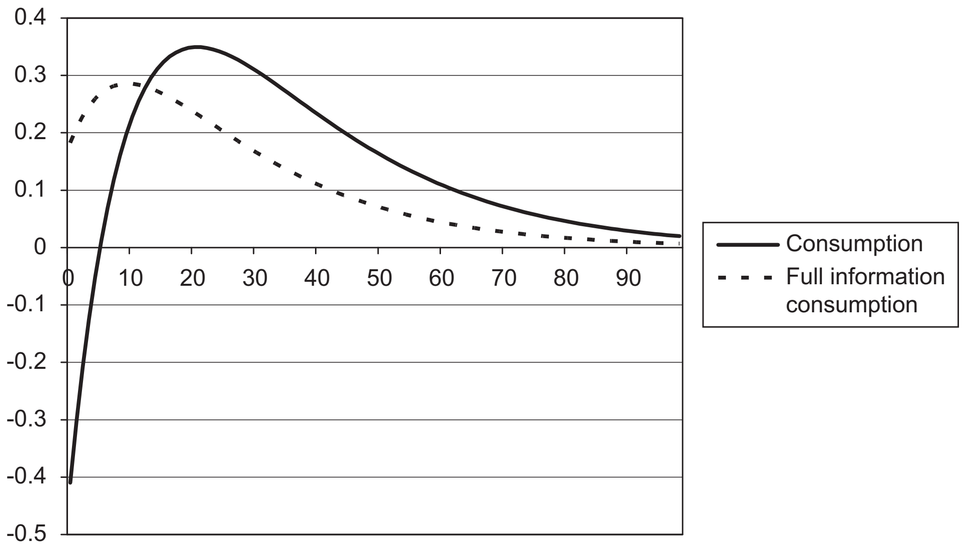

Fig. 1. Response of consumption to a 1% positive innovation to aggregate productivity.

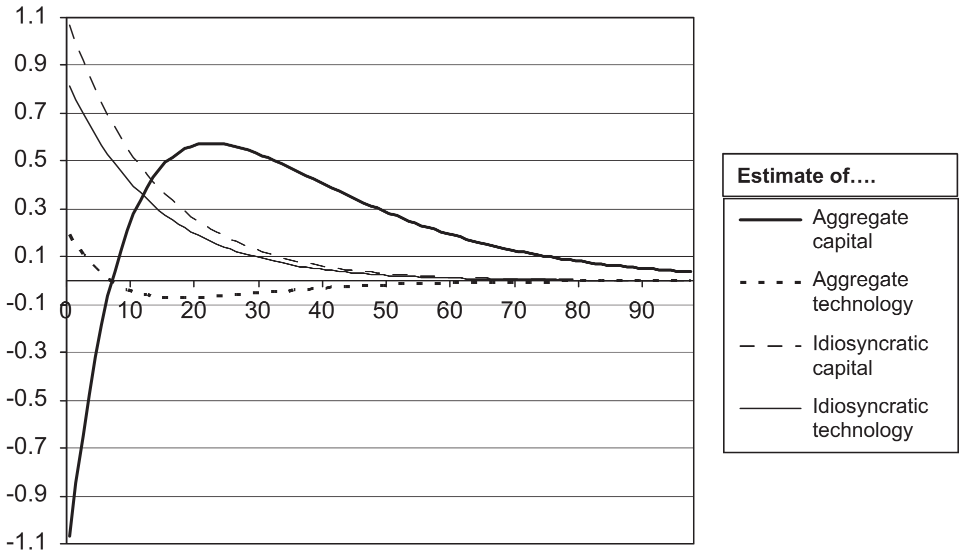

Fig. 2. Response of state estimates to a 1% positive innovation to aggregate productivity.

Innovations to the observed variables can occur either because of structural innovations, or because households' past state estimates were incorrect. For example, if there is a positive innovation to the return, this could either be caused by a positive aggregate technology shock, or because households had overestimated aggregate capital in the past. This is the basis for Corollary 1: households interpret what is actually good news about aggregate productivity as bad news about aggregate capital. This implies that technology must have also been overestimated in the past[11] which, given the persistence of technology, means that the sign of the response of the estimate of aggregate technology is ambiguous. In the base case the response is positive but quite small, around 20% of the size of the true shock. If technology is slightly more persistent than in the base case ($\phi_a = 0.93$), the two effects perfectly offset and estimates of aggregate technology never change.

The response of the state estimates are shown in Fig. 2. The predominantly bad news about the aggregate economy is offset by good news on both the idiosyncratic component of capital, $k_t^s - k_t$, since k_t^s itself is observed, and idiosyncratic

[11] To see this recall that estimates must be consistent with the information set so $r_{t-1} = E_{t-1}^s r_{t-1}$. Again assuming fixed labour for clarity, this implies $a_{t-1} - E_{t-1}^s a_{t-1} = k_{t-1} - E_{t-1}^s k_{t-1}$. If the household now believes $E_{t-1}^s k_{t-1}$ was too high this implies $E_{t-1}^s a_{t-1}$ was also too high. Since technology is persistent this will reduce the response of $E_t^s a_t$.

Table 1
Impact effect of aggregate technology shock on aggregate consumption: sensitivity to properties of idiosyncratic shock.

ϕ_z	σ_z/σ_a					
	∞	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

On the *x*-axis is the ratio of the standard deviation of the idiosyncratic shock to that of the aggregate shock, with the limiting heterogenous and homogenous cases at the left and right-hand sides. On the *y*-axis is the persistence of the idiosyncratic shock. The base case is shown in bold.

technology. But the pure permanent income response to estimates of idiosyncratic states is small, and so the overall response is dominated by the response to the estimate of aggregate capital.[12]

Households base their consumption decisions on estimates of the state variables. The accuracy of these estimates can be assessed by the covariance matrix of one-step ahead forecasts of the states. For the baseline calibration, under full information the quarterly standard deviations of one-step ahead forecast errors for aggregate technology and aggregate capital would be 0.7% and zero, respectively (since capital is pre-determined). In the base case with incomplete information the corresponding figures increase to 1.6% and 2.2%. It is striking that what seems to be a quite modest degree of uncertainty about the true value of the capital stock should be enough to cause such a significant change in the dynamics of the system, especially so, given that recent debates about the true size of the capital stock (see, for example, Hall, 2001 or the discussion of intangibles in Laitner and Stolyarov, 2003) have suggested measurement errors by statistical offices that are many orders of magnitude larger than this. The relative accuracy of households' estimates in our simple model suggests we may well be considerably understating the informational problem households face.

5. Robustness

This section discusses the robustness of the results to changes in the key parameters, and to the introduction of a noisy public signal and additional financial markets.

5.1. Sensitivities

The informational problem is due to the market-consistent information set being insufficient for the household to distinguish between aggregate and idiosyncratic productivity processes. The key parameters of the model are therefore the properties of these processes. Changing the other parameters does not affect the nature of the informational problem.

While there is strong evidence in the data that the idiosyncratic economy is much more volatile than the aggregate, an important question is the extent to which idiosyncratic shocks are observable. In informational terms, observing some part of idiosyncratic shocks is equivalent to their having a lower variance. Table 1 shows the sensitivity of the impact response of consumption to changes in the persistence and variance of the idiosyncratic shock.

Moving from left to right across the table, the degree of heterogeneity in the economy progressively falls, with the final column being the limiting homogenous case. As the relative standard deviation of the idiosyncratic shock decreases, the information problem becomes less acute, so the impact response of consumption becomes less negative.

Moving up the table, as the persistence of idiosyncratic technology rises, its unconditional variance rises so the informational problem becomes more acute with an increasingly negative impact on the response of consumption. However, there is a second offsetting effect. As the process becomes more persistent, an estimated innovation to idiosyncratic productivity has a greater effect on expected wealth, so, other things being equal, the response of consumption becomes less negative. These offsetting effects result in a non-monotonic relation between the persistence of idiosyncratic technology and the impact response of consumption.

5.2. A noisy public signal

The model presented so far has an informational structure which is internally consistent—in decentralized equilibrium the only source of information about aggregates are market prices—and this will never replicate the full information solution, so giving rise to welfare costs compared to full information. This could provide a rationale for the existence of

[12] This discussion is implicitly framed in "certainty-equivalent" terms, i.e. it assumes that the typical household's consumption only depends on estimates of the underlying states, W_t^s, whereas consumption also depends on the entire hierarchy of expectations. However, the *sign* of the actual consumption response is determined by the certainty-equivalent response.

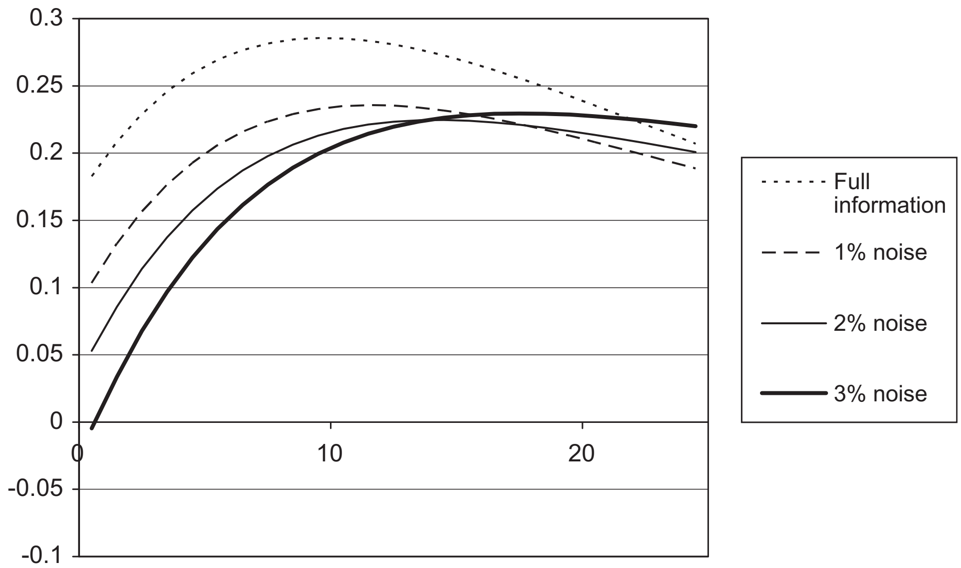

Fig. 3. Response of aggregate consumption to a 1% positive innovation to aggregate productivity with a noisy public signal of output.

other sources of information such as government statistical offices, etc. The impact of such additional sources of information can be analysed by extending the measurement vector (22) to include a public signal of output which differs from true output by a white-noise error.[13] Fig. 3 shows how this signal affects the response of aggregate consumption in the model with noise in the public signal with a standard deviation ranging from 1% to 3%.

Recall that without a public signal (Fig. 1) the impact response of consumption was negative. With a standard deviation of the noise in the output measure at the top of the range, the impact response becomes very close to zero. As the accuracy of the signal increases, the response of consumption approaches the full information case.

Although there is currently a lively debate on the empirical effect of technology shocks, see for example Christiano et al. (2003), there seems to be some agreement that a range of variables, including consumption, respond more sluggishly in the data than in a standard RBC model. Theoretical explanations for such sluggishness (for example Francis and Ramey, 2005) are usually couched in terms of nominal or real rigidities, or habit formation. The result of this section shows that informational imperfections can generate such a sluggish response of consumption without additional rigidities.

5.3. Financial markets

Our baseline model is one without financial markets. However the techniques of e.g. Lettau (2003) can be extended to price financial assets in this model. In contrast to the shadow prices normally derived under full information, in any model with incomplete markets and market-consistent information, these prices will have an informational content and so will change the equilibrium.

As an example, consider the case of a risk-free bond. Some algebra (see online Appendix (B5)) shows its return, r_{ft} is given by

$$r_{ft} = \mu' M X_t^{(1)} \tag{23}$$

where μ is the vector that relates the return on capital to the non-expectational aggregate states, M is as in (21), and $X_t^{(1)}$ is the first-order average expectation of the full state vector. So the risk-free rate reveals a linear combination of the hierarchy of expectations. However, in practice adding this to households' information sets barely changes the response in Fig. 1. The intuition for this is as follows. In the baseline calibration, with its large variance of the idiosyncratic shock, the wage conveys essentially no useful information about aggregates: the return on capital is effectively the only observable that gives information about aggregates, and all households know this to be the case. So even without a risk-free asset the hierarchy of average expectations of aggregate states is close to being common knowledge. Hence the risk-free rate conveys very little additional information, and the equilibrium is little changed. There are very similar results when a stock market is introduced.[14]

[13] How noisy are real-time estimates of output? Orphanides and van Norden (2002) attempt to quantify the extent of uncertainty by calculating the difference between real-time and final estimates. Their table 2 shows standard deviations of the difference ranging from 1% to 3% per quarter. However, they note that their method "…overestimate[s] the true reliability of the real time estimates since it ignores the estimation error in the final series", which given the issues involved in measuring output, is likely to be large but is by its nature unquantifiable.

[14] To price the stock assume a dividend process given by the one period return on capital as in Lettau (2003). The pricing equation is complicated by the need to allow for the hierarchy of expectations: the law of iterated expectations cannot simply be applied in solving forward. Instead substitute out for increasingly higher orders of average expectations (along very similar lines to the methodology used by Nimark, 2007). A closed form expression for the stock price, and hence the return, can be derived in terms of the full hierarchy, which for numerical solutions can again be truncated as outlined in Appendix B4. Precise details of the methodology are available from the authors.

What additional assets would change these results? The first type would be assets that help households insure against, and hence identify, idiosyncratic shocks. The more of these assets there are, the easier the filtering problem until in the limit with complete markets the economy collapses to the full information case (Radner, 1979). The second type are assets which give direct information about aggregate flows. For example, in our economy gross profits are a linear function of output, so observing these would be equivalent to observing output and would reveal full information. Since capital is owned by households, a stock market in this model would not reveal this information. Actual stock markets provide information about dividend flows, but in practice dividends and output are far from being perfectly correlated. We leave further examination of this issue to future work.

The relative unimportance of aggregate assets has an interesting parallel with the results of Athanasoulis and Shiller (2000), who find that, in a model of missing markets where new asset markets are added incrementally according to the magnitude of the resulting increase in welfare, the last asset to be added is the market portfolio itself, and that all preceding assets that are added are pure swaps, since these provide the crucial rule of risk pooling. In our framework, similarly, the assets that would provide the crucial information would be those that span the distribution of idiosyncratic shocks. Aggregate markets thus appear to play a marginal role both in terms of welfare and in terms of information.[15]

6. Conclusions

In marked contrast to the conclusions of recent research, we have shown that market incompleteness matters due to its informational consequences. It remains to be seen how robust this contrast will be to further modifications.

On the one hand it might be argued that we are overstating the informational implications of incomplete markets. Evidently markets are not so incomplete as in our model. While aggregate financial markets make very little difference, the existence of insurance and other risk-pooling markets would push these results closer to those under full information.

On the other hand, 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; symmetry across households; and a single aggregate endogenous state variable. Other models have more shocks and more states (for example Smets and Wouters, 2007, has seven shocks and four states) which will make the filtering problem of the household more complex, but may also have more sources of information. A striking feature of our model is how *well* households can estimate the capital stock, in stark contrast to the observed wide variation in estimates within the academic debate (e.g., Hall, 2001), which, it might be argued, reflect the much greater complexity of the true inference problem.

While our results are specific to the stochastic growth model, the principle of market-consistent information is much more general, and should be applied to any model of incomplete markets.

Appendix A. Supplementary data

Supplementary data associated with this article can be found in the online version at doi:10.1016/j.moneco.2009.12.005.

References

Angeletos, G.-M., La'O, J., 2009. Noisy business cycles. NBER Working Papers 14982.
Athanasoulis, S., Shiller, R., 2000. The significance of the market portfolio. The Review of Financial Studies 13 (2), 301–329.
Baxter, B., Graham, L., Wright, S., 2009. Invertible and non-invertible information sets in dynamic stochastic general equilibrium. Working Paper.
Campbell, J.Y., 1994. Inspecting the mechanism: an analytical approach to the stochastic growth model. Journal of Monetary Economics 33, 463–506.
Card, D., 1994. Intertemporal labor supply: an assessment. In: Sims, C. (Ed.). Advances in Econometrics, Sixth World Congress.
Christiano, L.J., Eichenbaum, M., Vigfusson, R., 2003. What happens after a technology shock? Working paper.
Fernandez-Villaverde, J., Rubio-Ramirez, F., Sargent, T., Watson, M., 2007. ABCs (and Ds) of Understanding VARs. American Economic Review, 1021–1026.
Francis, N., Ramey, V.A., 2005. Is the technology-driven real business cycle hypothesis dead? Shocks and aggregate fluctuations revisited. Journal of Monetary Economics 52, 1379–1399.
Guvenen, F., 2005. An empirical investigation of labor income processes. Working paper.
Guvenen, F., 2007. Learning your earning: are labor income shocks really very persistent? American Economic Review 97 (3), 687–712.
Hall, R.E., 1978. Stochastic implications of the life cycle-permanent income hypothesis: theory and evidence. Journal of Political Economy 86 (6), 971–987.
Hall, R.E., 2001. The stock market and capital accumulation. American Economic Review 91 (5), 1185–1202.
Hellwig, C., 2006. Monetary business cycle models:imperfect information. UCLA Working Paper.
King, R.G., 1983. Interest rates, aggregate information, and monetary policy. Journal of Monetary Economics 12 (2), 199–234.
Krusell, P., Smith Jr., A.A., 1998. Income and wealth heterogeneity in the macroeconomy. Journal of Political Economy 106 (6), 867–896.
Laitner, J., Stolyarov, D., 2003. Technological change and the stock market. American Economic Review 93 (4), 1240–1267.
Lettau, M., 2003. Inspecting the mechanism: the determination of asset prices in the RBC model. The Economic Journal 113, 550–575.
Lippi, M., Reichlin, L., 1994. VAR analysis nonfundamental representations Blaschke matrices. Journal of Econometrics 63, 307–325.

[15] King (1983) introduces a nominal bond into a Lucas-style monetary economy (without capital) and finds it mitigates informational problems. This is consistent with the present paper since the baseline model already has a common aggregate signal in the return on capital. The point in this section is that *additional* aggregate markets have little effect in a model with capital.

Lorenzoni, G., 2010. A theory of demand shocks. American Economic Review, forthcoming.
Nimark, K., 2007. Dynamic higher order expectations. Working paper.
Nimark, K., 2008. Dynamic pricing and imperfect common knowledge. Journal of Monetary Economics 55 (2008), 365–382.
Orphanides, A., van Norden, S., 2002. The unreliability of output-gap estimates in real time. Review of Economics and Statistics 84 (4), 569–583.
Prescott, E.C., 1986. Theory ahead of business cycle measurement. Federal Reserve Bank of Minneapolis Quarterly Review 10 (4), 9–22.
Radner, R., 1979. Rational expectations equilibrium: generic existence and the information revealed by prices. Econometrica 47 (3), 655–678.
Smets, F., Wouters, R., 2007. Shocks and frictions in US business cycles: a Bayesian DSGE approach. American Economic Review 97 (3), 586–606.
Townsend, R.M., 1983. Forecasting the forecasts of others. Journal of Political Economy 91 (4), 546–588.
Woodford, M., 2003. Imperfect common knowledge and the effects of monetary policy. In: Aghion, P., Frydman, R., Stiglitz, J., Woodford, M. (Eds.), Knowledge, Information, and Expectations in Modern Macroeconomics: In Honour of Edmund S. Phelps. Princeton University Press, Princeton.