

Heterogenous Information Choice in General Equilibrium*

Tobias Broer[†]

Alexandre Kohlhas[‡]

Kurt Mitman[§]

Kathrin Schlafmann[¶]

PRELIMINARY - PLEASE DO NOT QUOTE

Abstract

We study heterogeneous information acquisition in the dynamic equilibrium of a standard neoclassical economy. Heterogeneity in wealth holdings and labor market productivity, as in the standard Krusell and Smith (1998) environment, implies natural heterogeneity in the incentives to acquire information about the current state of the economy that allows to predict future wages and capital returns. This is because the savings choices of agents with low and high resources are approximately unaffected by future prices, making them unwilling to pay even low prices of information processing. Importantly, the benefit of acquiring information is declining in the mass of informed agents, whose countercyclical savings choices make the economy less volatile and information less beneficial. We show how, for some values of information cost, this implies that an equilibrium does not exist in the representative agent economy, while it exists with heterogeneous agents.

Keywords: Heterogenous information, rational expectations, bounded rationality

*Support from the Ragnar Söderbergs stiftelse is gratefully acknowledged. Mitman thanks the European Research Council (ERC Starting Grant 759482) for financial support.

[†]Institute for International Economic Studies and CEPR. Email: tobias.broer@iies.su.se.

[‡]Institute for International Economic Studies. Email: alexandre.kohlhas@iies.su.se.

[§]Institute for International Economic Studies and CEPR. Email: kurt.mitman@iies.su.se.

[¶]Institute for International Economic Studies and CEPR. Email: kathrin.schlafmann@iies.su.se.

1 Introduction

Rational expectations equilibria, the benchmark of modern macroeconomics, require agents to correctly perceive the stochastic transition laws that govern the dynamics of an economy. With many agents and endogenous heterogeneity, this requires substantial information about the distribution of individual state variables and individual decision rules. The resulting curse of dimensionality that inhibits solving for such an equilibrium has often led researchers to specify perceived equilibrium dynamics in heterogeneous agent economies in terms of a finite number of moments (typically simply the mean) of the cross-sectional distribution of individual state variables. The motivation for this is that higher moments of the distribution in any case do not help to predict variables of interest much in a statistical sense. So the inaccuracy of such “boundedly rational” expectations that neglect higher moments should imply only small costs.

We quantify the utility loss implied by different degrees of “bounded rationality”, or the use of different information sets, in the standard Krusell and Smith (1998) (henceforth KS) environment. We find that the benefits of acquiring information strongly depend on the information that other households use to form expectations. When all agents use a simple productivity-dependent autoregression to forecast future wages and returns (the Krusell and Smith (1998) benchmark), an individual household may neglect all moments but the unconditional mean of the capital distribution at low utility cost (equivalent to less than 0.05 percent of lifetime consumption). When households only use the unconditional mean of the capital distribution to forecast future wages and returns, the aggregate capital stock becomes substantially more persistent, and thus fluctuates more strongly in equilibrium. This makes the utility loss for an individual household of ignoring the current capital stock in her forecast of future variables sizeable (up to 3.5 percent of lifetime consumption). The benefits of conditioning forecasts on the current state of aggregate productivity, or the variance of the capital distribution, are always small, independently of the law of motion that households use in their forecasts in equilibrium.

These results have two main implications. First, informational equilibria in economies with heterogeneous agents might be sensitive to the cost of information acquisition. Only when information comes at negligible costs is the KS benchmark an equilibrium. When such costs are substantial, in contrast, agents may prefer to form expectations equivalent to the

unconditional mean of wages and returns. At intermediate costs, no symmetric equilibrium may exist. Second, information choice is likely to be heterogeneous. This is, first, because the costs of information acquisition may differ across households, and, second, because we find that the utility loss of using less information to form expectations differs strongly between households with different labor market status and wealth.

1.1 Relation to the literature

[To be added.]

2 A Baseline Model

We consider a simple two-period ($t = 1, 2$) economy with a continuum of households of mass 1 and information choice. Household i receives an endowment y of the single consumption good in the economy in the first period. She can transfer resources between the two periods by saving a non-negative amount $k \geq 0$ in an asset that pays a return of R in $t = 2$. Households income in the second period is the sum of the return on this investment, and the wage w earned from supplying inelastically one unit of labor. Households ignore the precise values of R and w , but perceive them to be distributed according to a prior probability distribution $P(w, R)$ on a bounded support $\Theta = [\underline{w}, \bar{w}] \times [\underline{R}, \bar{R}]$.

At the beginning of period 1 households have the option to pay a utility cost κ to acquire information that perfectly reveals R and w . This allows households to attain higher lifetime utility by taking into account their return and period 2 income in their savings decisions in period 1.

The household's problem is thus

$$\max_{\mathbb{I}_{inf}, k} E[\log(y_i - k) + \beta \log(Rk + w)] - \kappa \times \mathbb{I}_{inf} \quad (1)$$

where E denotes the expectation with respect to the prior probability distribution and \mathbb{I}_{inf} takes the value one when an individual acquires information and zero otherwise.

Below we consider information choice both from an ex ante perspective, when individuals do not know their income y , and ex post, after knowing y . For this, denote $U_{inf}(y)$ and $U_{noinf}(y)$ as the ex post expected utilities from informed and uninformed savings choices after

observing y , and $E_y[U_{inf}(y)]$ and $E_y[U_{inf}(y)]$, respectively, their ex ante counterparts, where E_y denotes the expectation with respect to the distribution of y , which itself may be a random variable.

2.1 Heterogeneous incentives for information acquisition

This section considers how the utility loss from not acquiring information, denoted as $\mathbb{L}(y) = U_{noinf}(y) - U_{inf}(y)$, and thus the incentives to acquire information, change as a function of first period income y . For this, note that the Euler equations for savings k with and without information choice are, respectively

$$\frac{1}{y - b} = \beta R \frac{1}{w + Rb} \quad (2)$$

$$\frac{1}{y - b} = \beta E_{R,w} \left[R \frac{1}{w + Rb} \right] \quad (3)$$

where E_x denotes the mathematical expectation with respect to the distribution of x .

The following three propositions summarize the information acquisition across the distribution of current endowments y . Proposition 1 shows that income-poor households, who are constrained and thus cannot change their savings as a function of information about w and R , never acquire information. The same holds for households rich enough such that $\frac{\bar{w}}{y} \approx 0$, as shown in Proposition 2: their consumption is approximately unaffected by realisations of w and, with log-preferences, their savings are also independent of the interest rate. So they have no incentive to acquire costly information about R . Finally, according to Proposition 3, when information costs are not too high, there are middle-income households that acquire information. All proofs are in the appendix.

Proposition 1 *A household with $y \leq \underline{y} = \frac{w}{\beta R}$ does not acquire information for any $\kappa > 0$.*

Proposition 2 *For any $\kappa > 0$ there is a value $\hat{y} > \underline{y}$ such that households with $y \geq \hat{y}$ do not acquire information.*

Proposition 3 *There is a $\kappa > 0$ and a $\tilde{y} > \underline{y}$ such that households whose first period income equals \tilde{y} acquire information.*

Figure 1 illustrates Propositions 1 to 3 by depicting \mathbb{L} , the expected utility loss from foregoing information, transformed into percentage differences in permanent consumption, as

a function of first period income y , for $\beta = 0.99$ and a uniform measure of $w \in [1, 10]$ and $R \in [0.9, 1.1]$.




Figure 1: The figure depicts the relative utility loss from not acquiring information in the simple two-period model for $\beta = 0.99$ and a uniform measure of $w \in [1, 10]$ and $R \in [0.9, 1.1]$

2.2 Equilibrium information choice

The previous subsection took the uncertainty about wages and returns in the absence of information, and thus the incentives to acquire it, as exogenous. In equilibrium, however, the distribution of future wages and returns is a function of agents current and prior decisions, including the information choice itself. To illustrate this circular relationship between information choice and the dynamic behavior of the economy that, in turn, determines the value of information, in this section we study the general equilibrium of the two-period economy described in the previous section. For this, we add two assumptions. First, we endogenise wages and returns through a standard competitive firm sector. Specifically, we assume that households save in capital that is rented out to competitive firms that use it, together with

the unit-supply of labor, in a standard Cobb-Douglas production technology

$$Y = K^\alpha L^{1-\alpha} \tag{4}$$

where K and $L = 1$ are, respectively aggregate capital and labor in period 2, and Y denotes aggregate output in that period. Markets for labor, capital, and consumption goods are competitive, so R and w are determined by the usual conditions, and thus a function of aggregate savings in capital K

$$w = (1 - \alpha)K^\alpha \tag{5}$$

$$R = \alpha K^{\alpha-1}. \tag{6}$$

This implies that information about second period prices is embedded in information about the aggregate savings K , as determined by individual savings choices in the first period.

In this environment, knowledge about the the distribution of period 1 endowments y , together with equilibrium savings choices, allows agents to perfectly predict second period prices w and R . The second assumption thus concerns the distribution of period 1 endowments and what agents know about it. Propositions 1 to 3 suggest that information choice divides the support of first-period income y into three ranges (with no information acquired for small and large values, and information acquisition in the middle). We choose the simplest distributional structure that captures this, together with uncertainty about aggregate savings choice, and thus future capital. For this, we assume that y takes on three values $y \in \{\epsilon, 1, \bar{y}\}$, where $\epsilon > 0$ and \bar{y} are small and large numbers respectively. We choose this support such that households at $y = \epsilon$ are always constrained to save 0 (and thus have no incentives to acquire information) and that period two wages are a negligible fraction of \bar{y} . So individuals with first period income \bar{y} always save a constant amount and have no incentive to acquire information.

To introduce a role for acquiring information about aggregate savings choices, we assume that households ex ante only know $\pi(1)$, the mass at $y = 1$. In the absence of information acquisition they ignore, however, how the remaining $1 - \pi(1)$ individuals are distributed between the extremes \underline{y} and \bar{y} . Specifically, there are two states of the world that differ only in the distribution of y . Agents attach a 50-50 prior probability to there being a high or low number of individuals at the upper bound \bar{y} , which translates to a 50-50 probability of a high

and low capital stock in period 2, respectively. They can, however, pay the cost κ to receive information about the true mass of agents at both point on the extremes. This allows them to accurately predict aggregate savings, and thus period 2 prices.

In the following we look at equilibria in this economy for different specifications of $\pi(1)$, the mass of agents at middle income, specifying ϵ and \bar{y} such that the aggregate period one endowment in both states remains constant.

Equilibrium

An equilibrium in this economy are price functions $w(K)$ and $R(K)$ in period 2, decision rules for information acquisition $I(y) \in \{0, 1\}$ and savings with $(k(y, K))$ and without acquisition of information $(k(y))$ in period 1, such that decision rules maximise expected utility for every agent, factor prices are given by (5) and (6) and period 2 capital equals the sum of savings.

Lemma 1 *The assumptions about the distribution of y imply that the equilibrium distributions of w and R fulfill the assumptions of Section 2.1. Specifically, w and R take values on finite, positive intervals $[\underline{w}, \bar{w}]$ and $[\underline{R}, \bar{R}]$, respectively.*

Proposition 4 *Non-existence of a pure-strategy equilibrium*

For each $\pi(1) \in (0, 1)$ there exists a cost $\kappa > 0$ such that no-pure strategy equilibrium exists.

The intuition behind proposition 4 is straightforward: Since optimal savings are strictly declining in w and increasing in R , informed unconstrained savings choices decline with aggregate capital in the second period ($k_2(y, K) < 0$). The dispersion of capital across the two states is thus strictly larger in an economy where no agent acquires information than in the economy where all middle-income agents are informed. Since the gain of acquiring information increases in the dispersion of second period prices, this implies that the gain of information acquisition is strictly smaller in the economy where all middle-income agents acquire information relative to that where nobody acquires it. There is thus a value of κ that makes middle-income agents want to acquire information when none of their peers do, but not when all of them do, implying that a pure-strategy equilibrium may not exist.

2.3 A numerical illustration

This section illustrates how information choice interacts with the stochastic equilibrium in the economy and how both are determined by the distribution of agents across period 1 income.

For this, we look at an example parameterization where $\beta = 0.99$ and $\alpha = 0.4$. Figure 2 depicts the standard deviation of capital in period 2 (left panel) and the ex ante expected consumption equivalent loss of not acquiring information (right panel) as a function of the mass of households at middle income $\pi(1)$ (along the bottom axis) when these middle-income households either acquire (dashed lines), or do not acquire information (solid lines).

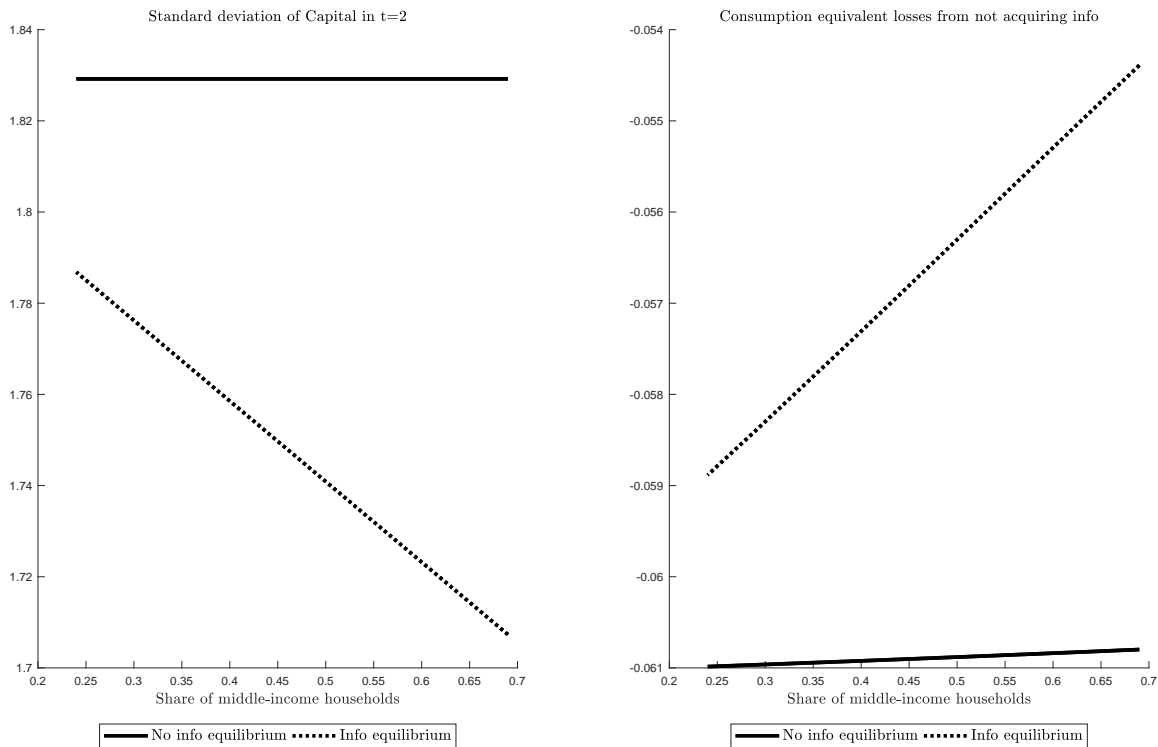


Figure 2: The figure depicts the standard deviation of capital in period 2 (left panel) and the ex ante expected consumption equivalent loss of not acquiring information (right panel) as a function of the mass of households at middle income $\pi(1)$ (along the bottom axis) when these middle-income households either acquire (dashed lines), or do not acquire information (solid lines).

When middle-income households acquire information, their savings are high when they know K will be low, and vice versa. This makes the second period capital stock less volatile than in the absence of information, when they save a constant amount that is approximately independent of their mass. This implies that the capital stock is less volatile when middle-income agents acquire information (as indicated by the difference between the dotted and solid lines in the left panel of Figure 2), and that the difference is increasing in the mass of middle




Figure 3: The figure divides the combinations of κ and $\pi(1)$ in Figure 2 into three regions, labeled by the type of equilibrium they imply.

income agents (as indicated by the downward slope of the dotted line). The utility loss from not acquiring information is therefore higher in the no-info equilibrium, when the capital stock is more volatile (in the right panel). Importantly, this difference is again increasing in the mass of middle-income agents $\pi(1)$: while the benefit of acquiring information is approximately independent of $\pi(1)$ in the no-info equilibrium (where nobody acquires information), the benefit of acquiring information in the info-equilibrium falls as $\pi(1)$ rises and more households save countercyclically (the dotted line in Figure 2 rises towards 0).

Figure 3 shows how this implies that there is range of values for the information cost κ where there is no pure-strategy equilibrium: costs are too high to make info acquisition worthwhile when all others acquire information, but too low for a households to forgo information when all others also do. This range of values is increasing in $\pi(1)$, the mass of middle-income households with an information choice. This implies, first, that when $\pi(1) = 1$,

and households thus have identical incomes, there may be no representative-agent equilibrium (as mixed-strategies imply different savings choices and thus a non-degenerate wealth distribution). Also, there is trivially a range of values κ such that there is no pure-strategy equilibrium when the mass of middle-income agents is high, while there is a unique pure-strategy equilibrium with information choice when the economy is more heterogeneous, in the sense that the mass of middle-income agents is lower. Heterogeneity, in the form of either income differences or random differences in strategies, is thus crucial for equilibrium existence.

3 The quantitative model

3.1 The general environment

The economic environment is a version of the KS economy with unemployment benefits, as studied in Den Haan et al. (2010) (and other articles in the same issue). Specifically, the economy consists of a continuum of ex ante identical households of unit mass. Households experience labor market shocks ϵ that make them transit from employment ($\epsilon = 1$) to unemployment ($\epsilon = 0$). When employed, a household earns wage w_t , when unemployed she receives unemployment benefits μw_t . The only asset in the economy is physical capital, whose net return equals $r_t - \delta$, the rental rate net of depreciation, and is equal for all households. Financial markets are thus incomplete and households can smooth consumption only through their choice of capital k_{t+1}^i by saving and dissaving, subject to a no-borrowing limit ($k_t^i > 0$). The household thus solves the problem

$$\max_{\{c_t^i, k_{t+1}^i\}_{t=0}^{\infty}} E_{\Omega_{i,t}} \sum_{t=0}^{\infty} \beta^t \frac{(c_t^i)^{1-\gamma} - 1}{1-\gamma} \quad (7)$$

$$s.t. \quad c_t^i + k_{t+1}^i = r_t k_t^i + (1 - \tau_t) \bar{l} \epsilon_t^i + \mu(1 - \epsilon_t^1) w_t + (1 - \delta) k_t^i \quad (8)$$

$$k_{t+1}^i \geq 0 \quad (9)$$

where $E_{\Omega_{i,t}}$ is the mathematical expectation conditional on information set $\Omega_{i,t}$ of household i in period t , c_t^i is household i 's consumption, \bar{l} her time endowment and τ a proportional tax on labor income.

Competitive firms rent capital and hire labor to produce the single output good in the

economy using a standard Cobb-Douglas production technology

$$Y_t = z_t K_t^\alpha (\bar{l}L_t)^{1-\alpha} \quad (10)$$

Markets for labor, capital, and consumption goods are competitive, so factor prices are given by

$$w_t = z_t(1 - \alpha) \left(\frac{K_t}{\bar{l}L_t} \right)^\alpha \quad (11)$$

$$r_t = z_t \alpha \left(\frac{K_t}{\bar{l}L_t} \right)^{\alpha-1} \quad (12)$$

where z_t is an exogenous process for aggregate productivity that takes on two values $1 - \Delta_z, 1 + \Delta_z$.

The only role of the government in this economy is to run a balanced-budget unemployment insurance scheme, such that

$$\tau_t = \frac{\mu u_t}{\bar{l}L_t} \quad (13)$$

where L_t and $u_t = 1 - L_t$ are, respectively, the employment and unemployment rates in the economy.

There are two exogenous sources of uncertainty in the economy: aggregate productivity z_t and individual employment status ϵ_t^i . Both follow a joint markov process, chosen such that the unemployment rate is a function only of z_t , and thus only takes on two values u_b, u_g with $u_b > u_g$.

3.2 Information-conditional General Equilibrium

In this section, we define a recursive competitive general equilibrium conditional on household information sets. Household decision rules and value function take as arguments their individual state variables k_t^i, ϵ_t^i , and the current state of the aggregate economy. Note that household decisions and cash-on-hand depend on aggregate capital through the wage rate and the return on investment, but are not directly affected by the distribution of capital. Given the assumption of incomplete markets, however, decision rules on individual capital holdings do not aggregate to a transition law for aggregate capital. In order to forecast future wages

and returns accurately, households thus have to take into account the joint distribution of capital and employment status in the economy, which we denote as Γ . To form expectations further into the future, moreover, households use a law of motion $\Gamma' = H(\Gamma, z, z')$, where H allows for dependence of Γ' on the realization of productivity z' through wages and returns.

The dynamic program of a typical household i is therefore

$$V(k, \epsilon; \Gamma, z) = \max_{c, k'} E_{\Omega} \frac{c^{1-\gamma} - 1}{1-\gamma} \beta V(k', \epsilon'; \Gamma', z') \quad (14)$$

$$s.t. \quad c + k' = r_t k + (1 - \tau) \bar{l} \epsilon + \mu(1 - \epsilon^1) w + (1 - \delta) k \quad (15)$$

$$k' \geq 0 \quad (16)$$

The joint distribution Γ is a high-dimensional object. To avoid the curse of dimensionality in their solution based on numerical dynamic programming, KS assume households to be *boundedly rational*, and perceive current and future prices to only depend on the first I moments of the joint distribution. They then use an iterative procedure to derive an approximate equilibrium. We assume a similar form of bounded rationality by conditioning the expectation operator in (14) on an arbitrary set of variables Ω that describes the state of the economy, and by assuming that households perceive Ω to follow a law of motion G . For now, we assume Ω to be the same across households and constant through time. Later, we will allow it to differ across households and time. Trivially, whenever Ω includes the current aggregate productivity state Ω and the first I moments of the distribution of capital, our approach is identical to that in KS.

To calculate an equilibrium, we choose an iterative procedure similar to that in KS: i) Choose Ω . ii) Postulate a law of motion G . iii) solve the consumer's problem conditional on Ω and G . iv) Using the resulting decision rules, simulate a large number of households for a large number of periods. v) From this simulation, calculate time series for the conditioning variables Ω , and estimate a new law of motion G' . vi) Compare this to G used in ii). If the two are different, update the guess for G and start again; if the two are sufficiently similar, stop.

Our approach differs from that in KS only insofar Ω may or may not include the aggregate productivity state or the mean of the capital distribution (in other words, we also consider $I = 0$), and that we do not necessarily require the resulting law of motion G to be an

Table I: Benchmark parameters

| | β | γ | α | δ | \bar{l} | μ | Δ_z |
|--------|---------|----------|----------|----------|-----------|-------|------------|
| Values | 0.99 | 1 | 0.36 | 0.025 | 1/0.9 | 0.15 | 0.01 |

Table II: Transition probabilities

| | $1 - \Delta_{z,0}$ | $1 - \Delta_{z,1}$ | $1 + \Delta_{z,0}$ | $1 + \Delta_{z,1}$ |
|--------------------|--------------------|--------------------|--------------------|--------------------|
| $1 - \Delta_{z,0}$ | 0.525 | 0.35 | 0.03125 | 0.09375 |
| $1 - \Delta_{z,1}$ | 0.038889 | 0.836111 | 0.002083 | 0.122917 |
| $1 + \Delta_{z,0}$ | 0.09375 | 0.03125 | 0.291667 | 0.583333 |
| $1 + \Delta_{z,1}$ | 0.009115 | 0.115885 | 0.024306 | 0.850694 |

accurate description of the dynamics of Ω . Hence the label “information-conditional general equilibrium”. After we have compared the properties of general equilibria conditional on Ω in the next section, we will consider several ways of determining equilibrium sets Ω chosen by households that may differ across individuals and time.

3.3 Parameter choice

As our benchmark we use exactly the same parameters as in Den Haan et al. (2010), given in table I and II.

4 The utility benefit of information

We are interested in the individual incentives to acquire information in KS-type economies. For this, we first solve for an equilibrium where all households use an information set $\tilde{\Omega}$. We then calculate the costs of “deviations”, where an individual uses an information set Ω that differs from the information set $\tilde{\Omega}$ used by all other households. To do this, we estimate laws of motion $G(\Omega)$ for the endogenous elements of this alternative information set Ω using the panel of individual variables generated in step iv) of the last iteration round of our algorithm. We then calculate decision rules for an individual who uses Ω to forecast the future. Finally, we calculate the expected utility from using those decision rules, evaluating expectations using a most comprehensive set Ω_{max} and its associated law of motion that are likely to allow accurate forecasts of the future. Finally, we calculate a relative utility loss of using Ω rather than the

most comprehensive Ω_{max} . Evidently, the resulting cost measure for deviations is conditional on the equilibrium information set $\tilde{\Omega}$, which may or may not coincide with Ω_{max} . In practice, we use the following formular:

$$CE_{\Omega}(k, \epsilon; \Omega_{max}) = (1 - \beta) [(1 - \gamma) (V_{\Omega}(k, \epsilon; \Omega_{max}) - V_{\Omega_{max}}(k, \epsilon; \Omega_{max})) + 1]^{\frac{1}{1-\gamma}} \quad (17)$$

where $\Omega \subset \Omega_{max}$, $V_{\Omega}(k, \epsilon; \Omega_{max})$ equals the expected discounted utility households with capital k and labor market status ϵ expect when they use the information set Ω and the aggregate state of the economy is described by particular values of the elements in Ω_{max} . Note that all terms in (17) also depend on the equilibrium law of motion $\tilde{\Omega}$, suppressed for simplicity. The utility cost of using only a subset of the potential information about the state of the economy Ω is thus the expected utility loss relative to a comprehensive set Ω_{max} , transformed into units of permanent consumption. Note that the incentives to acquire information may depend on both the individual state variables k, ϵ and the aggregate state of the economy, as given by specific values of the state variables (or their summary measures) in the comprehensive information set Ω_{max} . Moreover, as we will see, incentives also differ between the information sets $\tilde{\Omega}$ on which the equilibrium is conditioned.

4.1 Incentives for information acquisition in the KS economy

This section presents results for the benchmark economy in KS. In other words, we consider $\tilde{\Omega} = \{z, \bar{k}_t\}$, such that households' information sets contain aggregate productivity and the conditional mean of the capital distribution, and G takes the form of two simple autoregressions of \bar{k}' on \bar{k} conditional on z .

4.2 Incentives for foregoing information

We first look at the utility losses households suffer from having less information than contained in the KS $\{z, \bar{k}_t\}$. For this, we consider three different specifications of the alternative information set Ω : first, the empty set $\Omega = \emptyset$, implying that agents do not condition their expectations on anything, and therefore use only average transitions and the unconditional mean of capital in their forecasts for the future. Second, $\Omega = \{z_t\}$, implying that households can condition their expectations on the current aggregate productivity state, and use

productivity-specific unconditional means of k' in their forecast for the aggregate capital stock. And finally, $\Omega = \{\bar{k}_t\}$, such that agents know the current mean of the capital distribution but not the realisation of productivity z . This implies that they forecast the mean of capital in the next period using an unconditional AR(1) process, and use transition probabilities that do not condition on z to form expectations.

Figure 4 presents the utility cost of using less information in the KS economy, where $\tilde{\Omega} = \Omega_{max} = \{z, \bar{k}\}$, for different levels of individual wealth along the bottom axis, and the four different combinations of the aggregate productivity and individual employment states in different panels, evaluated at the median of the distribution of aggregate capital. As a reference, the red lines indicate the average cross-sectional distributions of individual capital in periods where the aggregate capital stock is within 1 percentage point of the median, conditional on the aggregate productivity state.

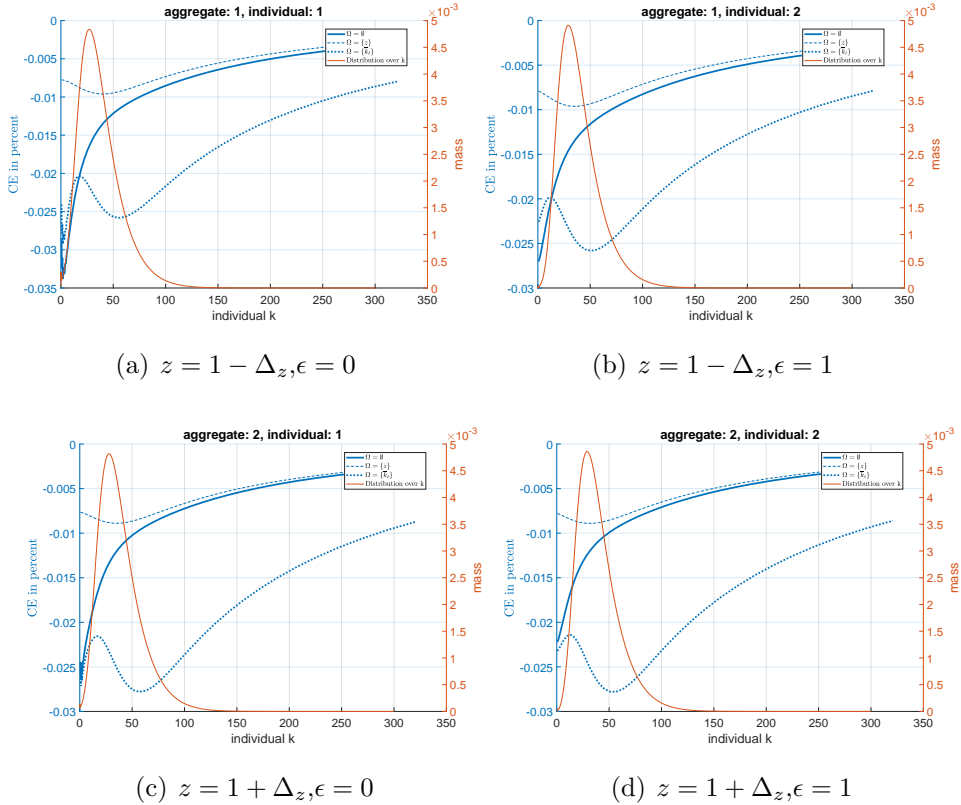


Figure 4: The figure presents $CE_{\Omega}(k, \epsilon; \Omega_{max})$, where $\Omega_{max} = \{z, k\}$ is the information set of KS, and Ω takes three specifications: first, the period-specific cross-sectional mean of k unconditional on z (implying \bar{k} is forecast using an unconditional autoregression for k , solid lines), z only (implying forecasts equal to the mean of k across simulations, conditional on aggregate productivity, dashed lines), and the empty set (implying forecasts of aggregate capital equal its unconditional mean, dotted line). All losses are evaluated at the median of the distribution of aggregate capital conditional on productivity.

Figure 4 shows that there is substantial heterogeneity in the utility loss from using less information in forecasting the future, but that the magnitude of losses is small for all values of the aggregate and individual states, and all specifications of Ω .

[Results to be discussed further.]

4.3 Incentives for adding information

Do households have incentives to acquire more information in the KS benchmark economy? KS show that the gains from adding additional moments to their benchmark improves the in-sample fit of the predictive regression for the capital stock in the future by only a small amount. They also argue that the welfare gains are “vanishingly small” (p. 878). We confirm this result: The maximum welfare gains from increasing the information set to also contain the variance, i.e. $\Omega = \{z_t, \bar{k}_t, var(k)_t\}$, are less than one thousandth of a percent of permanent consumption (where the maximum is taken over all individual and aggregate states).

4.4 Incentives for information acquisition in a low-information version of the economy

The last section showed that losses from suboptimal expectations are small in the KS economy. Does this mean agents should be expected to use simpler rules in equilibrium? Not necessarily. Figure 5 presents the same information as in 4 when the equilibrium information set is empty, so $\tilde{\Omega} = \emptyset$ and agents only use unconditional expectations in their forecast of future variables, and when we evaluate losses at the lower 10th percentile of the distribution of aggregate capital. The losses from using simpler rules that do not contain the current mean of the capital distribution are two orders of magnitudes larger than those contained in Figure 4. Figure 6 shows the reason for this: the dynamics of the aggregate capital stock strongly depend on the information set that the equilibrium conditions on. Specifically, when agents only use unconditional means to forecast the future, the aggregate capital stock is substantially more persistent than in the KS benchmark. This makes the distribution of the capital stock substantially more dispersed around its mean, and thus increases the errors from using the unconditional mean as an approximation for tomorrow’s capital stock.

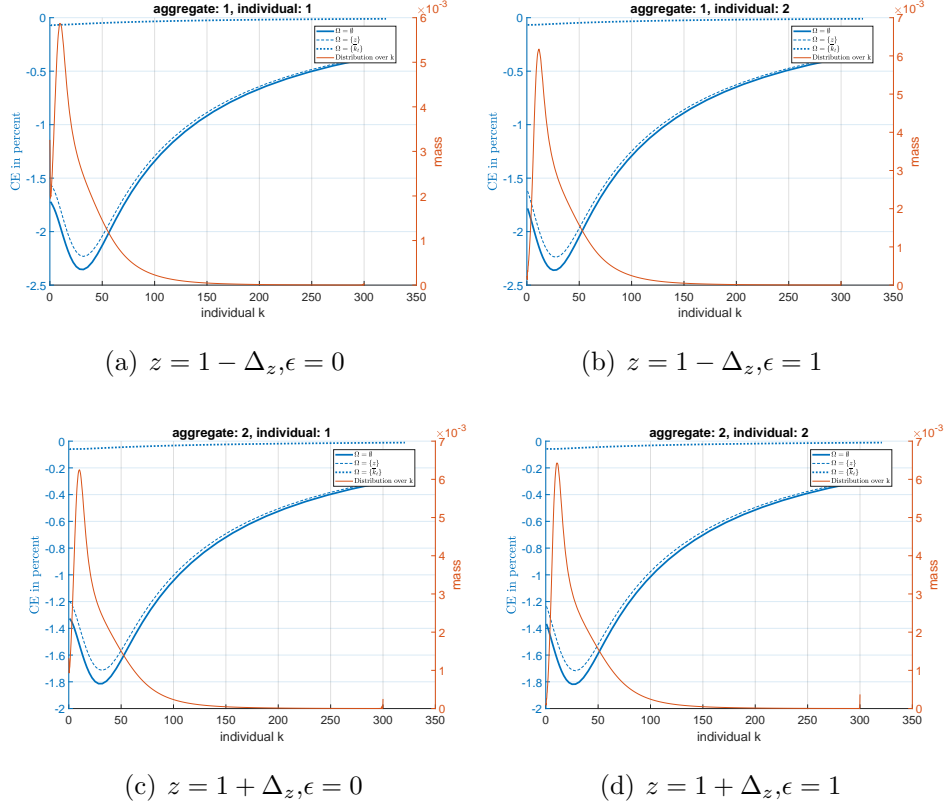


Figure 5: The figure presents the measure of utility losses $CE_{\Omega}(k, \epsilon; \Omega_{max})$ in a ‘low-information economy’ where $\tilde{\Omega} = \emptyset$ such that agents only use unconditional averages in forecasting the future. Ω takes three specifications: first, the period-specific cross-sectional mean of k unconditional on z (implying \bar{k} is forecast using an unconditional autoregression for k , solid lines), z only (implying forecasts equal to the mean of k across simulations, conditional on aggregate productivity, dashed lines), and the empty set (implying forecasts of aggregate capital equal its unconditional mean, dotted line). All losses are evaluated at the 10th percentile of the distribution of aggregate capital conditional on productivity.

Figure 6: Aggregate time series of capital




Table III: Expected losses in percent CE

| | $\tilde{\Omega} = \{z, \bar{k}_t\}$ | $\tilde{\Omega} = \emptyset$ | $\tilde{\Omega} = \{z\}$ | $\tilde{\Omega} = \{\bar{k}_t\}$ |
|--------------------------|-------------------------------------|------------------------------|--------------------------|----------------------------------|
| $\Omega = \emptyset$ | -0.0136 | -0.7050 | -0.7664 | -0.0048 |
| $\Omega = \{z\}$ | -0.0096 | -0.6841 | -0.7451 | -0.0017 |
| $\Omega = \{\bar{k}_t\}$ | -0.0270 | -0.0370 | -0.0378 | -0.0121 |

Table IV: Maximum losses in percent CE

| | $\tilde{\Omega} = \{z, \bar{k}_t\}$ | $\tilde{\Omega} = \emptyset$ | $\tilde{\Omega} = \{z\}$ | $\tilde{\Omega} = \{\bar{k}_t\}$ |
|--------------------------|-------------------------------------|------------------------------|--------------------------|----------------------------------|
| $\Omega = \emptyset$ | -0.0468 | -3.4892 | -3.1275 | -0.0280 |
| $\Omega = \{z\}$ | -0.0170 | -3.3257 | -2.9770 | -0.0024 |
| $\Omega = \{\bar{k}_t\}$ | -0.0466 | -0.1092 | -0.1419 | -0.0269 |

4.5 Summary measures of utility losses

This section summarizes the utility losses from using less than the full information set Ω_{max} using two measures: average and maximum expected losses, where both the average and maximum are taken over the whole ergodic distribution of individual and aggregate states.¹ Table III identifies the expected loss from using the three smaller information sets (corresponding to different rows) in four different economies that differ in the information that agents use in equilibrium (corresponding to different columns). For example, column 1 of table III shows the average losses in the KS economy ($\tilde{\Omega} = \{z, \bar{k}_t\}$), when agents use only unconditional means ($\Omega = \emptyset$, first row), conditional means ($\Omega = \{z_t\}$, second row), or an unconditional autoregression for the mean capital stock ($\Omega = \{\bar{k}_t\}$, third row) to forecast future variables. As suggested by figure 4, the expected losses are less than 0.05 percent of permanent consumption equivalents. The same is true when agents ignore the level of current productivity in equilibrium ($\tilde{\Omega} = \{\bar{k}_t\}$, column 4). Expected losses from using rules that ignore the current capital stock are two orders of magnitude larger, however, whenever the equilibrium information set does not include the current mean capital stock (columns two and three).

5 Equilibrium information choice

5.1 Symmetric information choice with ex ante criteria

[To be added.]

¹In practice, since the maximum losses in some cases we consider increase strongly at low capital values, we restrict ourselves to the 2.5-97.5 percentile range.

5.2 Heterogeneous time-varying information choice

[To be added.]

6 Appendix

Proof of Proposition 1: $y \leq \underline{y}$ implies that $U'(y) > \max_{R,w} [RU'(w)] \geq E_{R,w}[RU'(w)]$. So the household would not choose a positive k for any value of R and w , and its choice would therefore be unchanged by information.

Proof of Proposition 2: The strategy of the proof is to bind the loss of not acquiring information $\mathbb{L} = U_{noinf}(y) - U_{inf}(y)$ from above. For this, we first compute an upper bound for the difference in savings choices consistent with (2) and (3), respectively: we divide the state space into two sets of states where informed savings are either greater or not greater than the uninformed saving choice \hat{b} . The difference in savings in the first (second) set of states is bounded by those consistent with a prior degenerate at \underline{R} and \bar{w} (\bar{R} , \underline{w}), while the truth equals \bar{R} and \underline{w} (\underline{R} and \bar{w}). The respective mass of both states is bound above by one. Moreover, for every state, we can bound the losses of a wrong savings function from above by evaluating utilities at $\underline{w} = 0$. This yields a utility loss \mathbb{L} that is not smaller than a negative number that converges to 0 as y increases:

$$\mathbb{L} \geq \log\left(\frac{1}{1+\beta}\left(y + \frac{\bar{w}}{\underline{R}}\right)\right) + \beta \log\left(\frac{\bar{R}\beta}{1+\beta}\left(y - \frac{\bar{w}}{\underline{R}}\right)\right) \quad (18)$$

$$- \log\left(\frac{1}{1+\beta}(y)\right) - \beta \log\left(\frac{\bar{R}\beta}{1+\beta}(y)\right) \quad (19)$$

$$+ \log\left(\frac{1}{1+\beta}(y)\right) + \beta \log\left(\frac{\underline{R}\beta}{1+\beta}y + \bar{w}\right) \quad (20)$$

$$- \log\left(\frac{1}{1+\beta}\left(y + \frac{\bar{w}}{\underline{R}}\right)\right) - \beta \log\left(\frac{\underline{R}\beta}{1+\beta}\left(y + \frac{\bar{w}}{\underline{R}}\right)\right) \quad (21)$$

$$= \log\left(1 + \frac{\bar{w}}{\underline{R}y}\right) + \beta \log\left(1 - \frac{\bar{w}}{\underline{R}y}\right) + \log\left(1 - \frac{\frac{\bar{w}}{\underline{R}y}}{1 + \frac{\bar{w}}{\underline{R}y}}\right) + \beta \log\left(1 + \frac{\frac{\bar{w}}{\underline{R}y}}{\beta\left(1 + \frac{\bar{w}}{\underline{R}y}\right)}\right) \quad (22)$$

$$< 0 \quad (23)$$

where we choose y large enough for the non-negativity constraint on k to be slack, (18) is the utility from uninformed savings consistent with a degenerate prior at $w = \bar{w}$, $R = \underline{R}$ when the truth is $w = \underline{w} = 0$, $R = \bar{R}$, (19) is the utility from informed savings with $w = \underline{w} = 0$, $R = \bar{R}$, (20) is the utility from uninformed savings consistent with a degenerate prior at $w = \underline{w} = 0$, $R = \bar{R}$ when the truth is $w = \bar{w}$, $R = \underline{R}$, and (21) is the utility from informed savings with $w = \bar{w}$, $R = \underline{R}$. Clearly, as y increases, the resulting lower bound on \mathbb{L} converges to 0

monotonically from below. So for any $\kappa > 0$ there is a \hat{y} such that $\kappa > -\mathbb{L}$ for any $y > \hat{y}$.

Proof of Proposition 3: Choose some $\tilde{b} > 0$ and consider $\tilde{y} = \tilde{b} + \frac{1}{E[\frac{\beta R}{w+Rb}]}$ such that \tilde{b} solves the Euler equation without information acquisition for an individual with period one income $y = \tilde{y}$. Note that in any state of nature $\{R, w\} \in \Theta$ the utility implied by informed savings choices is never smaller than that implied by the uninformed savings choice \tilde{b} . To see how it is strictly greater in some states, note that the partial derivatives of b with respect to w and R implied by equation (2) are strictly non-zero when b is interior. Whenever the distribution $P(Y, R)$ is not degenerate, there exist thus $\{\hat{w}, \hat{R}\} \in \Theta$ with $P(\{\hat{w}, \hat{R}\}) > 0$, such that $b(\hat{w}, \hat{R}) - \tilde{b} = \epsilon$ with $abs(\epsilon) > 0$ when the individual receives information $\{\hat{w}, \hat{R}\}$. The strict concavity of period utility then implies that perturbing the optimal saving by a non-zero amount ϵ strictly reduces utility in state $\{\hat{w}, \hat{R}\}$. There is thus a κ small enough such that the individual optimally acquires information.

Proof of Proposition 4:

The proof has two steps. The first shows how the value of acquiring information is decreasing in the mass of informed agents. The second shows how this implies that for every $\pi(1)$, there is a nonempty range of values for κ such that no pure-strategy equilibrium exists.

It follows from (2), (5) and (6) that $k_2(y, K) < 0$. The linearity of the implied savings rule in (2) combined with that more information improves welfare $U_{inf}(y) - U_{noninf}(y) > 0$ then allows us to use Proposition 1 in Hellwig and Veldkamp (2009) to show the value of acquiring information $U_{inf}(y) - U_{noninf}(y)$ is strictly decreasing in the mass of middle income types that acquire information.

Clearly, for any mass of middle-income individuals $\pi(1)$, at $\kappa = 0$ all middle-income individuals choose to buy information, while for high enough κ none do. Now suppose that there was a pure-strategy equilibrium for all values of κ . This would imply that, for all $\pi(1)$, there is a cutoff value $\kappa(\pi(1))$ where $U_{inf}(y) - U_{noninf}(y) = \kappa$, and around which an infinitesimal increase in κ makes all middle-income individuals change from acquiring to not acquiring information. Since the net benefit of acquiring information $U_{inf}(y) - U_{noninf}(y)$ is strictly higher in the no-information equilibrium, however, this cannot be the case. In other words, for all $\pi(1)$, there must be a range of values for κ for which there are only mixed-strategy equilibria.

References

DEN HAAN, W. J., K. L. JUDD, AND M. JUILLARD (2010): “Computational suite of models with heterogeneous agents: Incomplete markets and aggregate uncertainty,” *Journal of Economic Dynamics and Control*, 34, 1 – 3, computational Suite of Models with Heterogeneous Agents: Incomplete Markets and Aggregate Uncertainty.

HELLWIG, C. AND L. VELDKAMP (2009): “Knowing What Others Know: Coordination Motives in Information Acquisition,” *Review of Economic Studies*, 76, 223–251.

KRUSELL, P. AND A. SMITH (1998): “Income and Wealth Heterogeneity in the Macroeconomy.” *Journal of Political Economy*, 106.