Heterogenous Information Choice in General Equilibrium*

Tobias Broer†

Alexandre N. Kohlhas‡

Kurt Mitman§

Kathrin Schlafmann¶

August 2020

Abstract

Motivated by the systematic differences in macroeconomic expectations across US households, we study optimal information choice in neoclassical economies with heterogeneous agents and incomplete markets. Incentives to acquire information about the current state of the economy differ strongly by wealth level and employment status. Because some households lose very little from uninformed choices, standard rational expectations equilibria are typically not robust to small costs of information. A calibrated version of our model matches key expectational patterns in US micro data. Both wealth inequality and business cycle volatility are substantially larger than with full information. Moreover, policies have an additional transmission channel through their effect on information choice. For example, by decreasing information acquisition, a wealth tax increases both aggregate volatility and inequality.

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.

†Paris School of Economics, 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.

¶Copenhagen Business School, Institute for International Economic Studies, Danish Finance Institute and CEPR. Email: ksc.fi@cbs.dk.

1 Introduction

Full information rational expectations equilibria, the benchmark of modern macroeconomics, require agents to correctly perceive the stochastic transition laws that govern the dynamics of an economy. In complex environments with many agents this requires substantial information about the distribution of individual state variables and decision rules that may be difficult, or costly, to acquire and process. At the same time, some agents may not benefit much from the improved prediction of future income and prices that accurate information about the current state of the economy allows. For example, households with little financial wealth may have little interest in predicting future returns. This suggests that we should expect incentives for information acquisition, and thus expectations, to be heterogeneous across agents. At the same time, these incentives may also change with macroeconomic conditions and policies.

In this paper, we first show how household expectations are indeed strongly heterogeneous in US micro data. According to the Survey of Consumer Expectations (SCE), both the location and dispersion of reported probability distributions for key macroeconomic variables differ substantially among survey respondents. Importantly, this heterogeneity partly reflects significant differences across gender, educational attainment, labor market participation and wealth. Motivated by this evidence, and the complexity of rational expectations equilibria, we study optimal information choice in the dynamic general equilibrium of an otherwise neoclassical heterogeneous-agent economy. We first show how heterogeneity in wealth holdings and employment status, as in the standard Krusell and Smith (1998) (henceforth KS) environment, naturally implies heterogeneity in the incentives to acquire information. Low-wealth households, who value current consumption highly, decide not to pay even small monetary costs of information acquisition. Similarly, depending on their risk aversion, a constant savings rule may have low utility costs for agents with substantial levels of financial wealth and relatively little labor income. Households that are not poor but whose lifetime income is dominated by future wages, in contrast, value information about current conditions highly, as it allows them to better predict their employment prospects and thus to make better savings decisions. We show how in equilibrium even small information costs therefore imply substantial heterogeneity in the degree of information across agents.

Importantly, relative to the full information rational expectations benchmark, such limited information acquisition substantially changes the equilibrium properties of the economy, and the effects of policies. Deviations from steady state are larger and more persistent, and wealth and consumption inequality is substantially increased. At the same time, increased volatility of aggregate conditions raises the benefits of information about the current state of the economy. While this dampens the effect of limited information in our environment, we still obtain large heterogeneity in information acquisition with a substantial fraction of agents optimally not acquiring information in any given period. Importantly, this endogeneity of information acquisition gives policies additional effects. We show this with the example of a wealth tax, which reduces average information acquisition by reducing average wealth levels.

These results have two main implications. First, rational expectations equilibria are typically not robust to small costs of information acquisition. This is because, with heterogenous agents, some of them typically lose very little from uninformed choices. In addition, a representative agent equilibrium may not exist, as individual benefits of information are low when the level of information in an economy is high, and vice versa. Second, with heterogeneous information choice and small information costs, aggregate fluctuations are amplified, inequality is substantially increased, and policies have additional effects on fluctuations and inequality by changing information acquisition.

We first illustrate the heterogeneity in incentives for information acquisition, and their relationship with aggregate equilibrium dynamics, in a simple two-period model. With log-preferenes, households naturally split into three groups according to their first-period resources: The first group are poor households, for whom costs of information acquisition outweigh the limited benefits they can obtain from information. The second group consists of those households who are rich enough for consumption to be approximately unaffected by future wages (and for whom the income and substitution effects of interest rates cancel each other out). They will not pay any utility cost of information even if it perfectly reveals future wages and returns and therefore do not acquire information either. This result depends, however, on the assumption of log-preferences: with higher risk-aversion, there is a threshold of current resources beyond which households always acquire information. The final group consists of households with an intermediate level of current resources. They strictly benefit from information acquisition that improves savings decisions, and thus acquire information as long as costs are not too high.

Because savings of informed households are high when they expect low capital and wages in the second period (and vice versa), informed savings reduce the dispersion of future wages and interest rates, and thus lower the cost of uninformed savings decisions today. In other words, benefits of information are low when the fraction of information acquirers is high, and vice versa. We show how this implies a range of information costs for whom there is no equilibrium with homogeneous strategies. By implication, there may not exist a representative-agent equilibrium in an economy that consists of ex ante-identical agents. We show how these results extend from the two-period model to an infinite-horizon neoclassical incomplete markets economy with aggregate and idiosyncratic risk and once-and-for-all information choices, as in KS.

The simplified framework thus highlights the importance of endogenous heterogeneity for information choice, its inherently dynamic character (as household incentives change when aggregate or individual conditions evolve), and the importance of general-equilibrium feedback from aggregate dynamics to individual choices. This motivates us to study fully dynamic information choice in a standard incomplete-markets economy with idiosyncratic and aggregate risk, when there are small costs of information acquisition. Agents decide every period whether or not to update a dynamic prior about the current level of productivity and the aggregate capital stock by acquiring information.[1] We choose a constant monetary cost of information and a stochastic utility cost to target the mean and standard deviation of expectational errors in SCE data. This implies that households update their information on average once every six quarters. The probability of information acquisition (before i.i.d. utility shocks are realized) is strongly heterogeneous across different priors, wealth levels and employment states. Since the wealth distribution and priors are continuous, this solves the non-existence problem. Information acquisition probabilities are highest for the unemployed with positive but moderate asset holdings, as future wage and employment prospects (which are well predicted by the current state of the economy) are crucial for their savings choices but monetary costs not too onerous in utility terms. Households whose incomes are diversified between wages and financial

[1]We avoid the typical problem that arises when combining heterogeneous and endogenous information, namely that an infinite number of higher-order beliefs matters for the equilibrium allocation (see, for example, Townsend (1983)), by focusing on information acquisition about exogenous productivity, exploiting the fact that, in neoclassical heterogeneous-agent economies, knowledge of the sequence of productivity shocks allows extremely accurate predictions also about aggregate capital.

returns, in contrast, lose relatively little from uninformed consumption and savings choices. With risk aversion higher than 1, the gains of information acquisition then rise again as wealth increases towards the top of the distribution. As a result of this heterogeneity, expectational errors follow an inverse U-shape in wealth, as in the data. Importantly, wealth inequality is substantially amplified by limited information, with 90/10 percentile ratios about 30 percent higher. In addition, macroeconomic volatility is increased, with standard deviations of output and investment around 10 percent higher than with full information.

Apart from these effects of limited information on equilibrium features at a given set of parameters, we show how macroeconomics policies may have substantially different effects when one takes into account their effect on endogenous information acquisition. For this, we look at the example of a wealth tax that finances wasteful government spending. A 0.5 (1) percent p.a. wealth tax reduces information acquisition by between 15 (30) percent, as average information acquisition falls with the decrease in wealth accumulation. Importantly, while the implied fall in average aggregate capital is similar to that with full information, the impact of the wealth tax on aggregate volatility and cross-sectional inequality is substantially altered through the additional transmission channel via endogenous information acquisition. Specifically, while aggreate volatility is not much affected under full information, the standard deviations of output and capital in our benchmark economy increase by 3 and 9 percent, respectively, as lower information increases the dispersion of the aggregate capital stock. The effect of the tax on inequality differs even more strongly: a one percent tax reduces the 90/10 and 99/1 percentile ratios of the wealth distribution with full information (by 4 and 7 percent), but increases them by 2 percent with limited information. These results suggest that their effect on endogenous information acquisition may substantially alter the relative costs and benefits of macroeconomic policies.

1.1 Relation to the literature

Our analysis is related to three separate strands of research. First, we contribute to the large literature on heterogeneous-agent macroeconomics following KS.[2] Our contribution to this literature is, first, to show that such analyses are typically not robust to small costs of information, as some households have no or small incentives to make accurate forecasts. Second,

[2]See Krueger et al. (2016) for a recent survey.

we show that small costs of information acquisition imply a large degree of heterogeneity in information in equilibrium, with a substantial fraction of agents optimally not acquiring information in any given period. Relative to KS-type equilibria (where agents' perceived law of motion is required to fit the actual law of motion well in a statistical sense), aggregate dynamics are amplified and wealth and consumption heterogeneity is substantially increased.

Second, our work is related to the large literature on the macreconomic and financial effects of information choice. Prominent studies are, among others, Grossman and Stiglitz (1980), Mankiw and Reis (2002), Sims (2003), and Hellwig and Veldkamp (2009a). Specifically, our paper contributes to the literature that studies the general equilibrium effects of information choice (e.g. Maćkowiak and Wiederholt (2009), Maćkowiak and Wiederholt (2015), Angeletos and Lian (2018)). We emphasize the role of heterogeneity in income and wealth, and how that combines to shape households' information choices in a dynamic framework. The contribution of our paper is, in this context, to highlight the macroeconomic consequences of dynamic, heterogenous information choice.

Finally, our study contributes to the literature on dynamic macroeconomic models with heterogeneous expectation formation (see Branch and McGough (2018) for a survey). In particular, our work is similar in spirit to those studies that allow rationally heterogeneous expectations as a function of random cost shocks (as, for example, in Brock and Hommes (1997)), and that take a temporary equilibrium approach to aggregate heterogeneous individual behavior. Similar to that literature, we also highlight the strategic substitutability in information choice as a feature of the limited information equilibrium. In contrast, heterogeneous expectations in our framework arise from rational heterogeneous information choice as a function of heterogeneous wealth levels, employment status and priors about the current state of the economy, in an environment where decision rules do not aggregate. Moreover, we keep track of the whole, endogenous distribution of households across individual state variables over time, which allows us to highlight the circular relationship between wealth accumulation and information acquisition.

Section 2 summarizes key patterns of expectations in SCE data. Section 3 presents a model of dynamic information choice in an environment with aggregate and idiosyncratic risk and incomplete markets. Section 4 derives analytical results in a simplified, two-period version of the model. Section 5 presents the quantitative analysis. Section 6 concludes.

2 Heterogeneity of household expectations in US data

This section documents stylized facts about households' expectations of key macroeconomic variables. For this we use data from the Survey of Consumer Expectations (SCE), a monthly internet survey administered by the Federal Reserve Bank of New York that since 2013 asks a rotating panel of about 1300 US households for their expectations about a number of individual and macroeconomic variables. Importantly, the survey features probabilistic questions, which ask respondents to indicate their perceived probability distributions of some future macroeconomic outcomes, such as consumer and house price inflation.[3] We compare these household expectations to outcomes, and forecasts from the US Survey of Professional forecasters (SPF), the main quarterly survey of macroeconomic forecasts in the US.

We document three stylized facts about economic expectations of US households.

1. Household expectations are substantially less accurate and less precise than professional forecasts. Thus, median root-mean squared forecast errors for 12-month CPI inflation by households are more than twice as large as those by professionals (equal to 1.8 and 0.7 percentage points (pp), respectively). And a median interquartile range of *individual* forecast distributions of 2 in SCE data makes household forecasts substantially less precise than those by professionals (whose median interquartile range equals 0.6 pp).

2. Household expectations are substantially more heterogeneous than professional forecasts. For example, households differ strongly in their assessment of future labor market conditions: the stated probabilities of higher unemployment 12 months from the survey date have an average interquartile range of 30 percentage points (pp), compared to 20 pp for professional forecasts. The relative heterogeneity of household inflation forecasts is even higher, both in terms of their point estimate and their precision.[4] Moreover, the few other forecasts of macroeconomic variables elicited in the SCE (such as house price inflation) show very similar patterns.

3. Household expectations in the SCE are correlated with individual economic conditions.

[3]See Section 2 in the appendix for a description of the data.

[4]For 12-month CPI inflation, point forecasts have average interquartile ranges of 2.6 and 0.9 pp for households and professional forecasters respectively. And individual forecast precisions (as indicated by the interquartile ranges of individual forecast distributions) have a range of 1.8 and 0.3 pp, for households and professional respectively, between quartiles 1 and 3 of survey respondents.

Figure 1: Household expectations across the wealth distribution

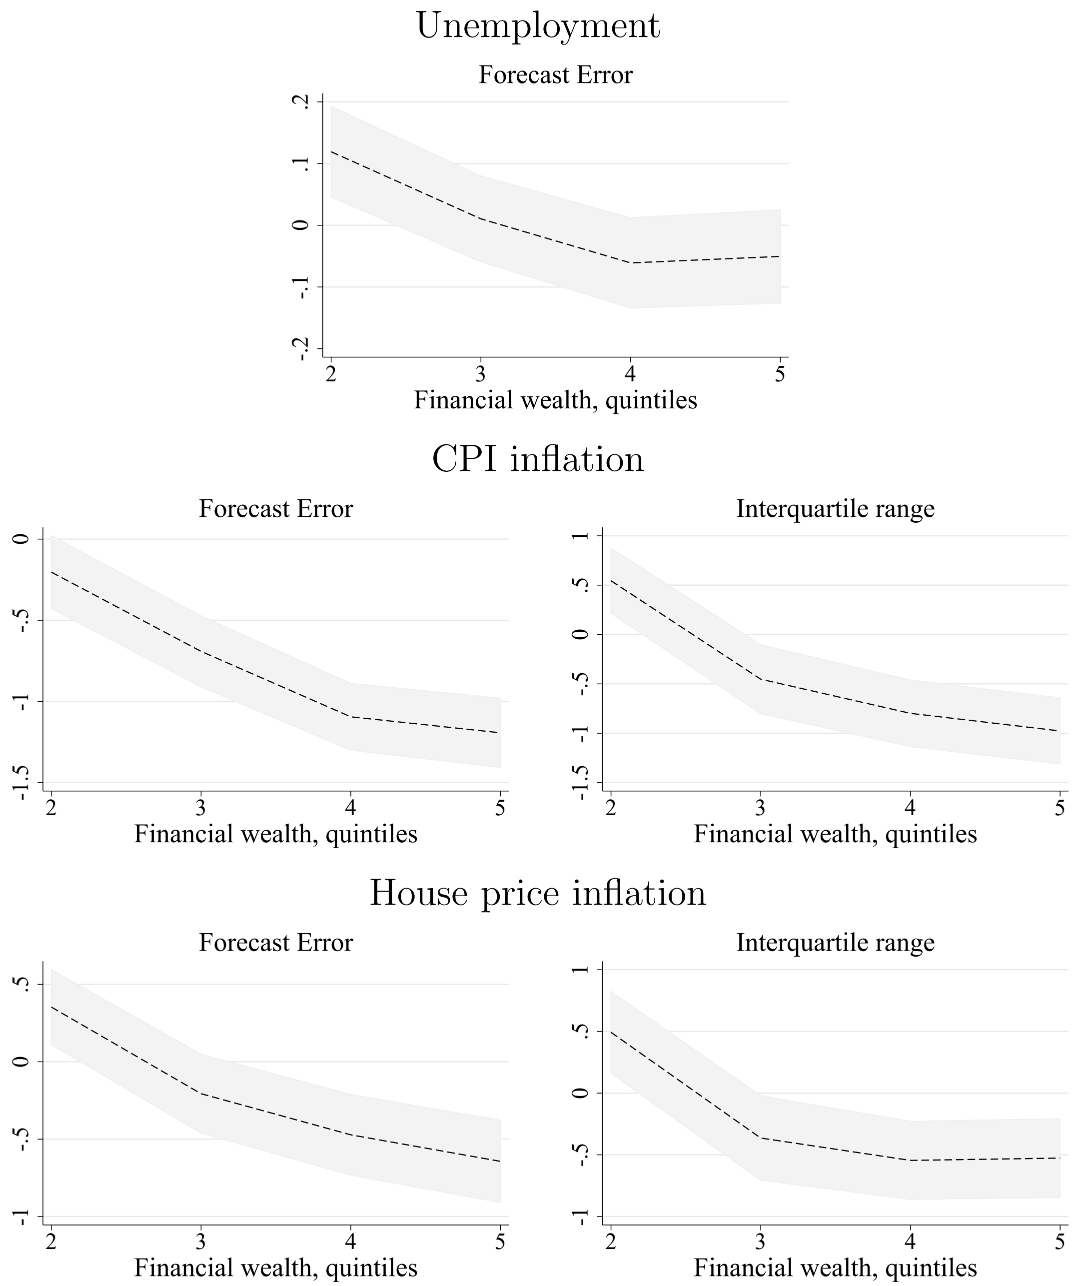

Based on the estimates of the random-effects models in Table VII, the Figure shows the indicator variables that take value one if a respondent reports wealth in the ith quintile of the overall wealth distribution, for $i = 2, ..., 5$, plus one-standard-error bands around them. The top row shows estimates for the errors in individual unemployment forecasts (elicited as the "percent chance that 12 months from now the unemployment rate in the U.S. will be higher than it is now") relative to the equivalent consensus forecast from professional forecasters (see Appendix 2 for detail). Row 2 presents results for root-mean-squared forecast errors (left panel) and individual forecast precision (as indicated by the interquartile range, right panel) for 12 month-ahead consumer price inflation. Row 3 shows equivalent results for house price (HP) inflation.

So observed heterogeneity is not just due to noise. In particular, households with higher wealth make on average smaller errors and have tighter forecast distributions. Moreover, as shown in Figure 1, there is evidence that this effect is not monotone but hump-shaped. Apart from the point-forecasts for inflation, all measures of forecast errors and forecast dispersion are higher for households in the second wealth quintile than for the wealth poorest households.[5] We also find signifcant differences across labor market participants and non parcipants, and households with different education and gender.[6]

The rest of the paper shows how heterogeneous optimal information acquisition in an otherwise standard neoclassical heterogeneous-agent economy explains some of these stylized facts about household expectations of macroeconomic variables.

3 A heterogeneous-agent economy with dynamic information choice

This section presents a standard neoclassical infinite-horizon economy where households can choose the information they receive about the current state of the economy every period. To allow for heterogeneity of incentives to acquire information across the wealth and income distribution, we assume that households face idiosyncratic unemployment risk which they cannot fully insure because financial markets are incomplete, as in KS.

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, indexed by i, whose preferences are described by a standard utility function

$$U = \sum_{t=0}^{\infty} \beta^t \frac{c_t^{1-\gamma} - 1}{1-\gamma}. \tag{1}$$

[5]The p-values of the corresponding coefficients are between 0.1 and 0.14.

[6]See Table VII in the appendix, which shows the results from panel regressions of moments of individual forecasts on a number of household characteristics (including a cubic function in age and dummies that correspond to the five quintiles of the wealth distribution in the sample) and year-month fixed effects.

Households experience idiosyncratic labor market shocks ϵ_t that make them transit from employment ($\epsilon_t = 1$) to unemployment ($\epsilon_t = 0$). A household earns wage w_t when employed, and receives unemployment benefits μw_t when unemployed. 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 trough their choice of capital k^i_{t+1} by saving and dissaving, subject to a no-borrowing limit ($k^i_t \geq 0$).

Competitive firms rent capital and hire labor in order to maximise profits from producing the single output good in the economy using a standard Cobb-Douglas production technology

$$Y_t = z_t K_t^\alpha (\bar{l} L_t)^{1-\alpha} \tag{2}$$

where K_t and L_t denote, respectively, capital and employment in period t, and $\bar{l}$ is the labor endowment. Markets for labor, capital, and consumption goods are competitive, so factor prices are given by

$$w_t = z_t(1-\alpha)\left(\frac{K_t}{\bar{l}L_t}\right)^\alpha \tag{3}$$

$$r_t = z_t \alpha \left(\frac{K_t}{\bar{l}L_t}\right)^{\alpha-1} \tag{4}$$

where z_t is an exogenous process for aggregate productivity.

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

$$\tau_t = \frac{\mu u_t}{\bar{l} L_t} \tag{5}$$

where τ_t is a tax rate on labor earning, L_t and $u_t = 1 - L_t$ are the employment and unemployment rates in the economy, respectively.

There are two exogenous sources of uncertainty in the economy: aggregate productivity $z_t \in \{Z_l, Z_h\}$ and individual employment status $\epsilon^i_t \in \{0, 1\}$. Both follow a joint markov process described by a four-by-four markov matrix Π with typical element $\Pi_{z',\epsilon'|z,\epsilon}$, 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

$u_b > u_g$.[7] We denote by Π^z the two-by-two markov matrix of (marginal) transition probabilities for z_t, with typical element $\Pi^z_{z'|z}$.

3.2 The informational setting

We assume that households take a dichotomous decision every period whether to acquire information about the current state of the economy or not. For this, we impose a "maximum" information set $\Omega_{max,t}$ that contains a small number of variables, or signals, that describe the state of the economy. Every period, households choose a subset $\Omega_{it} \subseteq \Omega_{max,t}$ of these signals.

Because savings decisions are a non-linear function of individual capital holdings, there is no aggregate law of motion for capital that households could use to forecast future wages and returns. The state of the economy thus comprises, in principle, the full (infinite-dimensional) joint distribution of capital and employment status. Motivated by the fact that, in KS-type economies, a law of motion that uses information about the level of productivity z_t and the current mean of the capital distribution $\bar{k}_{t+1}$ allows extremely accurate predictions of future prices, we take as a benchmark $\Omega_{max,t} = \{z_t, \bar{k}_t\}$. Consistent with rational expectations, we assume that agents use the equilibrium law of motion G to update their posterior distribution about future variables, conditional on their information.[8] Households have to pay both a monetary cost ν and a utility cost κ that are functions of the variables they choose to observe. Reflecting ample evidence that expectations are heterogeneous even for households with similar observable characteristics, we assume that κ is stochastic, but independently distributed over time.

Our assumptions are similar to KS, who assume that, apart from the current productivity level z, households use the first I moments of the current cross-sectional distribution of capital holdings to forecast future prices every period. Trivially, whenever Ω_{it} is constant over time and identical for all households, and includes the current aggregate productivity state z and the first I moments of the distribution of capital, our approach is identical to that in KS. In contrast to KS, however, we also consider information sets that do not include the aggregate productivity state or the mean of the capital distribution (in other words, we also consider

[7] See Krueger et al. (2016) for details on the restrictions this imposes for Π.

[8] This assumption renders our model tractable as it directly limits the number of signals that agents can choose to learn from. Without any assumption restricting the number of signals, agents would in principle use the entire distribution of individual state variables to form their (rational) expectation.

$I = 0$), and we allow them to differ across households and time as a function of idiosyncratic and aggregate state variables . Moreover, while we retain the requirement that, in equilibrium, households' perceived law of motion G capture accurately the conditional distribution of elements in Ω_{it}, we do not necessarily require G to describe their dynamics accurately in a statistical sense. In other words, we allow households to make non-negligible expectational errors in equilibrium.

3.3 Household problem and equilibrium

Every period, decisions are taken in two stages: In a first stage, after idiosyncratic employment shocks ϵ are realized, households decide whether to acquire information or not. Specifically, households choose a set of signals $\Omega_{it} \subseteq \Omega_{max,t}$ that contain information about the current level of productivity z_t and / or mean capital $\bar{k}_t$, and form a posterior accordingly. Thus, their information set $\mathbb{I}_{it}$ accumulates as $\mathbb{I}_{it} = \{\Omega_{it}, \mathbb{I}_{it-1}\}$. In a second stage, conditional on their posterior, households make consumption and savings choices.

Stage 1

In the first stage households thus choose the information to acquire $\Omega \in \{\emptyset, \{z\}, \{\bar{k}\}, \{z, \bar{k}\}\}$:

$$V_1(y, \epsilon, \kappa; \Phi_1(z, \bar{k})) = \max_{\Omega} E_{\Phi_1} \left[V_2(y, \epsilon, \kappa, \Omega; \Phi_2(z, \bar{k})) \right] \tag{6}$$

$$\Phi_2(z, \bar{k}) = \Theta(\Omega, \Phi_1(z, \bar{k})) \tag{7}$$

where the expectation is taken with respect to the prior distribution about z and $\bar{k}$, denoted $\Phi_1(z, \bar{k})$. V_1 and V_2 are, respectively, the value functions before and after information acquisition. They take as arguments individual state variables (cash on hand y, her labor market status ϵ, and the realization of her utility cost κ) and the current information about aggregate conditions as summarized by, respectively, the prior and posterior distribution about z_t and $\bar{k}_t$, denoted $\Phi_1(z, \bar{k})$ and $\Phi_2(z, \bar{k})$. V_2 takes as an additional argument the acquired information Ω (that determines costs paid in the second stage). Finally, Θ in (7) updates the prior $\Phi_1(z, \bar{k})$ in line with that information choice.

Stage 2

In the second stage, households choose consumption and savings given $\Phi_2(z, \bar{k})$

$$V_2(y, \epsilon, \kappa, \Omega; \Phi_2(z, \bar{k})) = \max_{c,k'} \left[\frac{c^{1-\gamma} - 1}{1-\gamma} - \kappa(\Omega) + \beta E_{\Phi_2} V_1(y', \epsilon', \kappa'; \Phi_1(z', \bar{k}')) \right] \tag{8}$$

$$s.t. \quad c + k' + \nu(\Omega) = y \tag{9}$$

$$y' = r'k' + (1-\tau)\bar{l}\epsilon' w' + \mu(1-\epsilon')w' + (1-\delta)k' \tag{10}$$

$$\bar{k}' = G(z, \bar{k}, z') \tag{11}$$

$$c, k' \geq 0 \tag{12}$$

where the expectation in (8) is taken with respect to the posterior distribution $\Phi_2(z, \bar{k})$, and the monetary cost ν and utility cost κ are a function of information acquisition choice as given by the set Ω. (9) is the standard budget constraint, for cash on hand defined in (10) as the sum of the return on capital holdings (net of depreciation) and labor (or replacement) income. Households use the equilibrium law of motion (11) and the exogenous transition matrix Π^z to form a prior $\Phi_1(z', \bar{k}')$ about next period's productivity and mean capital from today's posterior $\Phi_2(z, \bar{k})$. Note that the true aggregate state of the economy is not an argument of the individual value functions.

General equilibrium

Given the informational assumptions above, the definition of a general equilibrium is similar to that in KS. In particular, a recursive competitive equilibrium in this economy is summarized by a law of motion G, the individual's value functions V_1 and V_2, the policy functions for consumption c, capital k' and information choice Ω, as well as pricing functions r and w such that the value and policy functions solve the household problem, r and w are given by (3) and (4), G is generated by exogenous transitions and individual policy functions, and the markets for capital and labor clear.

4 Analytical results in a simplified model

In this section we look at a simplified version of the economy, whose analytical tractability allows us to highlight the economic forces that determine heterogeneous information choice in equilibrium. For this, we collapse the future to one period $t = 2$, and consider information choice in $t = 1$, conditional on exogenous prior information and cash on hand. We proceed in two steps. Since the relevant variables for consumption and savings choices are future prices, in Section 4.1 we let households directly acquire information about second-period wages and returns conditional on their cash on hand. This allows us to highlight the heterogeneity of information choice across the wealth distribution in partial equilibrium.

In a second step, in Section 4.2, we look at information acquisition in general equilibrium, where period-two prices are determined by productivity shocks and aggregate capital, as determined by period-one savings. Because informed savings are "countercyclical" (high when the future capital stock is known to be low and vice versa), the ex-ante dispersion of period-two outcomes and the benefit of information acquisition declines with the fraction of information acquirers. We show how this implies that a pure strategy equilibrium may not exist.

4.1 Partial equilibrium information choice

Consider first the problem of a household with cash on hand y, and assume, for simplicity, that she has the option to purchase at utility cost κ accurate information about period-two wages w and returns R (whose time subscripts we drop in this section), before choosing period-one consumption and savings. Households enter the first period with some non-degenerate prior distribution $\Phi_1(w, R)$ about w and R on some bounded support $\Psi = [\underline{w}, \overline{w}]x[\underline{R}, \overline{R}]$, for $\underline{w} < \overline{w}$ and $\underline{R} < \overline{R}$. The Euler equations for savings k with and without information choice are, respectively

$$(y - k)^{-\gamma} = \beta R(w + Rk)^{-\gamma} \tag{13}$$

$$(y - k)^{-\gamma} = \beta E_{\Phi_1} \left[R(w + Rk)^{-\gamma} \right] \tag{14}$$

Proposition 1 *Consider households in the simple two period model with an exogenous prior about second-period prices w, R.*

i. A household whose current income y is low enough to save 0 even at minimal discounted wage (such that $y \leq \underline{y} = \underline{w}(\beta\bar{R})^{-\frac{1}{\gamma}}$) does not acquire information for any information cost $\kappa > 0$.

ii. There is a strictly positive value of information cost $\kappa > 0$ and a value of income $\widetilde{y} > \underline{y}$ such that households whose first period income equals $\widetilde{y}$ acquire information.

iii.a *When $\gamma = 1$ (log-preferences), for any information cost $\kappa > 0$ there is a threshold value of income $\widehat{y} > \underline{y}$ such that households with income above the threshold ($y \geq \widehat{y}$) do not acquire information.*

iii.b *For every coefficient of risk aversion γ exceeding 1, there is a positive information cost $\kappa > 0$ and a threshold value of income $\overline{y}_\gamma$ such that households whose income exceeds that threshold $y \geq \bar{y}_\gamma$ acquire information.*

Part i. of proposition 1 shows that households whose income is low enough to save 0 regardless of the present discounted value of future wages never pay for information. With log-preferences, the same holds for households rich enough such that $\frac{\overline{w}}{y} \approx 0$, or non-participants in the labor market without wage income, as their savings are unaffected by wages and independent of R (iii.a). When risk-aversion is higher than 1, there is a threshold of current resources such that those with endowments beyond that level do acquire information (iii.b). When information costs are not too high, there are middle-income households that acquire information (ii.).

Figure 2 illustrates Proposition 1 numerically, by depicting $\mathbb{L}$, the expected utility loss from uninformed savings and consumption choice for different prior distributions, transformed into percentage differences in permanent consumption as a function of first period income y. With log-preferences (left-hand panel), losses follow an inverse U-shape pattern, and approach zero as the endowment y rises. Expected losses are lower when wages and interest rates are perceived to be less volatile, or to comove more strongly (making the present discounted value of future wage payments that determines informed savings less variable). With higher risk-aversion equal to 5 (right-hand panel), losses are higher and no more inverse-U shaped. Rather, losses have a local minimum at intermediate values of endowments y, where sources of future income are most "diversified" across wages and returns on savings, which is more powerful when wages and interest rates correlate negatively. As y rises further, losses converge to a strictly positive limit (as suggested by part iii.b of Proposition 1).

4.2 Equilibrium information choice

The expected losses in Figure 2 depend strongly on the individual prior distribution over wages and returns. In equilibrium, prices are given by (3) and (4), and thus functions of period-two productivity z (again dropping time subscripts), and the distribution of cash-on-hand y that together determine savings and thus capital K. In contrast to the full model, the two-period

Figure 2: Utility loss L in the two-period model

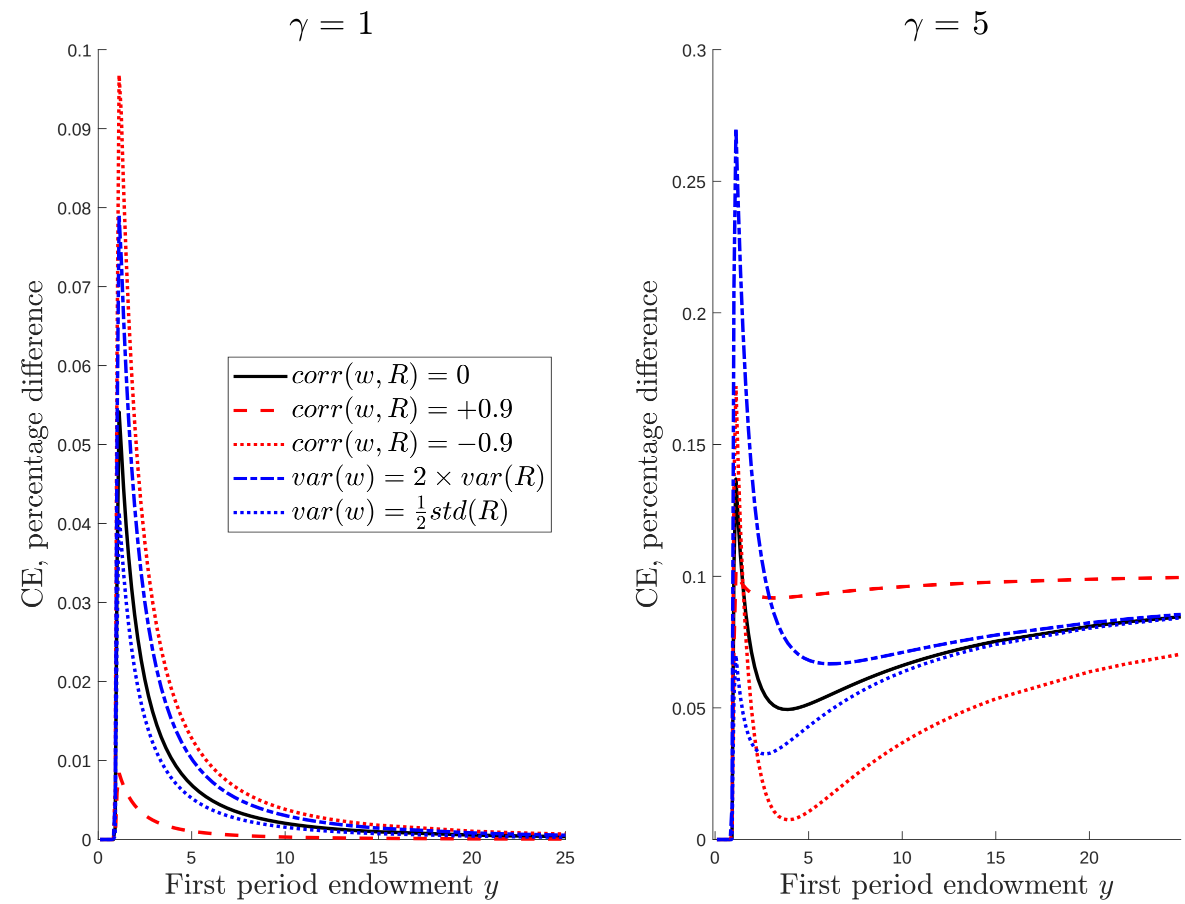

The figure depicts the relative utility loss from not acquiring information in the simple two-period model for $\beta = 0.99$ and a joint normal distribution of w and R with means of 1 and standard deviations of 5 percent. The left-hand panel considers the case of log-preferences ($\gamma = 1$), the right-hand panel that of risk-aversion equal to 5.

example allows us to study separately simple cases where households can directly acquire information about z and about (a conveniently chosen) distribution of y, with exogenous priors. For this, we concentrate on the case of log-preferences.

Proposition 2 (Information about z: Unique no-information equilibrium) *With log-preferences ($\gamma = 1$), when households know the current distribution of y, no household chooses to buy information about the period-two productivity level z.*

The proof follows immediately from substituting (3) and (4) into (14) with $\gamma = 1$, which implies that savings are independent of z. Households who know the distribution of y therefore correctly anticipate aggregate savings, and do not acquire information.

Individuals do have incentives, in contrast, to acquire information about the current distribution of y, which determines future wages and returns via aggregate savings. To illustrate this simply, we assume that y takes on three values $y \in \mathbb{Y} = \{\epsilon, 1, \overline{y}\}$ and denote the mass of agents as $\pi(x), x \in \{\epsilon, 1, \overline{y}\}$. We choose the 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 $\overline{y}$ (implying, with log-preferences, that information about R leaves savings unaffected, and is thus worthless). So only middle-income households make a non-trivial information choice. We choose a simple source of uncertainty about the endowment distribution: households know $\pi(1)$ and the aggregate endowment, but have a a 50-50 prior about how many of the remaining agents have cash on hand $y = \overline{y}$ (where ϵ adjusts to keep aggregate endowments constant). Because of the higher marginal propensity to save of the rich, this translates into a 50-50 prior over aggregate savings, and thus over future wages and returns.

Proposition 3 (Information about distribution of y: Non-existence) *With log-preferences ($\gamma = 1$), when households know the level of productivity z, there exists for each $\pi(1) \in (0, 1)$ a cost $\kappa > 0$ such that no homogeneous pure strategy equilibrium exists.*

Figure 3 illustrates Proposition 3. It reflects the strategic substitutibility of information choice: the savings of middle-income households are high when they know that the mass of rich agents, and thus aggregate savings, are low, and vice versa. This decreases the dispersion of capital across the two possible endowment distributions, and thus the value of information. The standard deviation of capital as perceived by the uninformed thus declines with the mass of informed middle-income agents $\pi(1)$ when costs are such that these acquire information (in the "info-equilibrium", the dotted line in the left panel of Figure 3). There is thus an increasing difference in the perceived dispersion of capital between the info- and "no-info" equilibrium (where costs are such that nobody acquires information, the solid line). The difference in utility losses between the two equilibria thus also increases with the mass of middle-income agents: while the benefit of acquiring information is approximately independent of $\pi(1)$ in the no-info equilibrium (the solid line in the right panel of Figure 3), the benefit of acquiring information in the info-equilibrium falls as $\pi(1)$ rises and more households save countercyclically (the

Figure 3: Equilibrium information choice in the two-period model

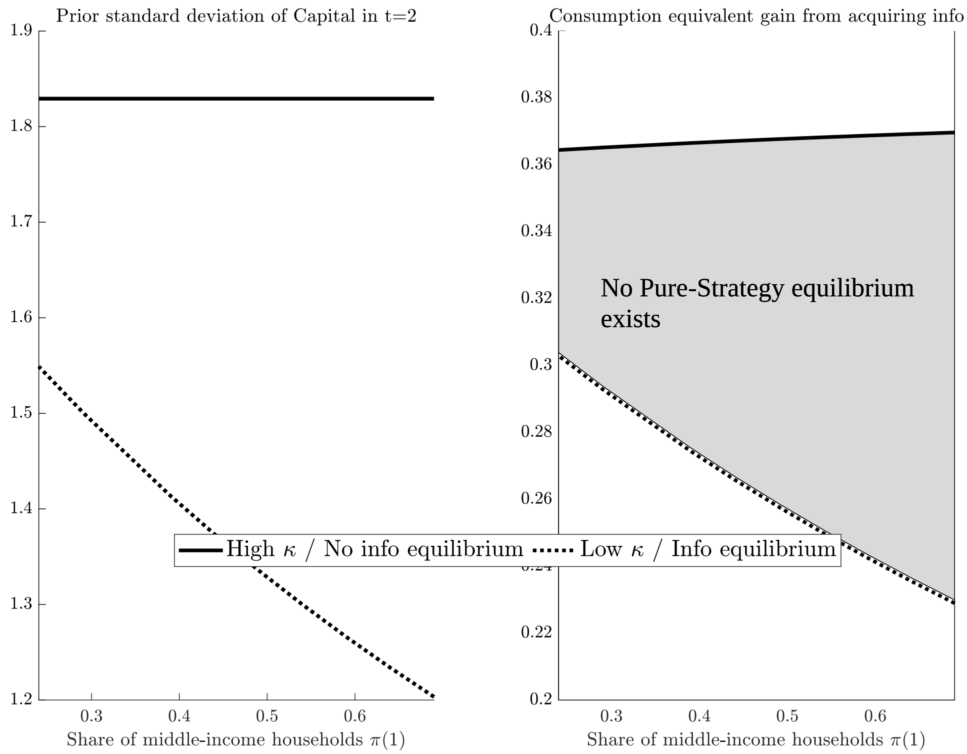

The figure depicts the equilibrium standard deviation of capital in period 2 (left panel) and the ex ante expected consumption equivalent loss of not acquiring information (right panel, in percent), both evaluated at the prior distribution of uninformed agents, as a function of the mass of households 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). The remaining mass of households is split between high-income households (whose income equals 6500 and whose mass is either $\pi^h(\overline{y}) = 0.08$ or $\pi^h(\overline{y}) = 0.06$ percent. For every $\pi(1)$ we choose ϵ to normalise aggregate endowments to 1 (checking that the resulting value is consistent with zero savings). Also, we use illustrative parameters $\beta = 0.99$ and $\alpha = 0.4$.

dotted line). In the simple example of this section, this implies that there is range of values for the information cost κ, increasing in $\pi(1)$, where there is no pure-strategy equilibrium (the gray area): for any individual middle-income household costs are too high to make information acquisition worthwhile when all other middle-income households acquire information, but too low for a household to forgo information when all others also do. Importantly, when $\pi(1) = 1$, there may therefore not be a representative agent equilibrium (where information and savings choices are the same for all middle-income households).

Corollary 1 *When $\pi = 1$, there is a range of values for κ such that there is no equilibrium with a representative middle-income agent.*

The simple model of this section illustrates the heterogeneous incentives to acquire information across the wealth distribution, and their interdependence in the population because of strategic substitutibility. The distribution of current resources was inevitably exogenous, and only had limited support (giving rise to the potential for non-existence). Moreover, most of the results concentrated on the case of log preferences. The rest of this paper analyses quantitatively the full model, with more general preferences and an endogenous continuous time-varying wealth distribution whose moments households forecast using a perceived law of motion.

5 Quantitative results

The simplified two-period version of our model allowed us to show analytically how heterogeneity in wealth and employment status translates to heterogeneous incentives for information acquisition, and how homogeneous-information equilibria may not exist (because aggregate volatility and gains from information tend to be high when average information levels are low, and vice versa). Appendix 7.5 shows that very similar results obtain in a quantitative analysis of our benchmark infinite-horizon economy with once-and-for-all information choice.[9] This suggests that we should allow information choices to vary across indidividuals that differ in wealth and employment status and across time (as idiosyncratic conditions change), in a general equilibrium environment that allows for feedback from aggregate dynamics to individual choices and vice versa. This section presents such an analysis of dynamic heterogeneous information choice in general equilibrium. The results show how information acquisition policies are indeed strongly heterogeneous in the cross-section, and less than a fifth of households acquire information every period. This substantially increases wealth inequality and aggregate volatility relative to a full-information version of the economy. And importantly, the endogenous information acquisition gives macroeconomic policies an extra transmission channel, whose importance we show using the example of a wealth tax.

[9] In particular, the losses from uninformed choices show a very similar pattern across the wealth distribution to that in Figure 2 (but are higher for unemployed households who benefit more from using information to predict future employment probabilities). And homogeneous-information equilibria again do not exist for intermediate costs of information, as, again, the volatility of output and of the capital stock are substantially higher when households do not acquire information about the aggregate state of the economy. In other words, the KS-equilibrium, where all households acquire information that allows statistically accurate forecasts of future prices, is not robust to small costs of information, as the utility cost of uninformed consumption and savings choices is small for many agents.

5.1 Priors and information sets

To study dynamic heteregeneous information choice in general equilibrium, we make two additional assumptions. The first allows us to circumvent the well-known "Problem of Infinite Regress of Expectations" (Townsend (1983)) that arises with endogenous signals. Specifically, we exploit a well-known feature of neoclassical, heterogeneous-agent economies with full information: a similarly accurate prediction of future prices to that using z_t and $\bar{k}_t$ results by using the sequence of shocks $\{z_s\}_{s=0}^{t-1}$ to form a prior about $\bar{k}_t$. In other words, while information about the current value of productivity z_t does not reduce losses much relative to no information as it does not help predicting the current capital stock (as suggested by the results in Table VIII in the previous section), knowledge about the sequence of past realizations does. We therefore consider the simple case in which agents decide each period whether or not to acquire information about the exogenous value of current productivity z_t. The function Θ in (7) thus simply updates the prior probability of being in a boom, denoted ϕ_1, to a posterior probability ϕ_2 equal to 1 or 0 whenever information is acquired.

Second, we assume that households attach a Dirac delta distribution to the observation of the mean capital stock, so that agents disregard uncertainty about $\bar{k}_t$ and have a point prior denoted $\hat{k}$. This makes the household problem tractable (while any continuous, or even discretized, prior distribution over $\bar{k}_t$ would add a high-dimensional state variable to the individual problem, and thus make it intractable). Based on a posterior $(\phi_2(z), \hat{k})$, the prior next period is thus simply

$$\phi_1(z')_= \phi_2(z)\Pi^z_{22} + (1 - \phi_2(z))\Pi^z_{12} \tag{15}$$

$$\hat{k}' = \phi_2(z)G(Z_h, \hat{k}) + (1 - \phi_2(z))G(Z_l, \hat{k}) \tag{16}$$

Section 7.4 in the Appendix presents this version of the households problem in detail.

5.2 Parameters and solution procedure

Table I summarizes the parameters we use in our quantitative analysis. We interpret a time period as a quarter. We choose a relative risk aversion parameter γ equal to 5. This choice is informed by the evidence that prediction errors are declining in wealth in US micro data (Figure 1) at medium and high wealth levels. We interpret this as incentives to acquire information

Table I: Benchmark parameters

	β	γ	α	δ	$\bar{l}$	μ	Z_l	Z_h
Values	0.99	5	0.36	0.025	1/0.9	0.40	0.984	1.0104

that are increasing in wealth, and thus choose γ higher than 1 (which would imply declining benefits of information, see Figure 2). We choose standard parameters for the discount factor β (0.99), the capital share α (0.36), and the depreciation rate δ (0.025). We calibrate the structure of aggregate and idiosyncratic uncertainty to capture key features of the dynamics of unemployment and job-finding rates in the post-world war II US economy, in the spirit of KS. Specifically, we specify transitions in aggregate productivity to capture good and bad times, defined as periods when unemployment is below and above trend, respectively.[10] The productivity then captures the difference in average US total factor productivity during the periods thus identified. The resulting persistence of good and bad times is 0.88 and 0.82, respectively, similar to that in KS (who use a symmetric persistence of 0.875). The resulting values for Z_l, Z_h are 0.9838 and 1.0104, respectively. The parameters governing individual transition probabilities are specificied to be similar to those observed in the US labor market. In particular, we choose an unemployment rate in booms and recessions equal to 6 and 10 percent, respectively. Job-finding rates are set such that unemployment spells are relatively short, as in US data, equal to 55 and 45 percent in booms and recessions, respectively. The remaining transition probabilities are then pinned down by the requirement that the unemployment rate depend only on current productivity, and reported in Table II. Finally, we normalise the labor endowment $\bar{l}$ to have unit labor supply in the bad aggregate state, and set the replacement rate μ equal to 0.4.

To calculate an equilibrium, we choose an iterative procedure similar to that in KS: i) Postulate a law of motion G. ii) solve the consumer's problem conditional on G. iii) Using the resulting decision rules, simulate a large number of households for a large number of periods. iv) From this simulation, calculate time series for z_t and $\bar{k}_t$, and estimate a new law of motion G'. v) Compare this to G postulated in i). If the two are different, update the guess for G and

[10]We use an hp filter with smoothing parameter 14400 to construct the trend in the unemployment rate at monthly frequency.

Table II: Transition probabilities

	$0\|Z_l$	$1\|Z_l$	$0\|Z_h$	$1\|Z_h$
$0\|Z_l$	0.5500	0.4500	0.5500	0.4500
$1\|Z_l$	0.0500	0.9500	0.0056	0.9944
$0\|Z_h$	0.4500	0.5500	0.4500	0.5500
$1\|Z_h$	0.0777	0.9223	0.0351	0.9649

start again; if the two are sufficiently similar, stop.

Information cost parameters

We choose information cost parameters that qualitatively capture key features of the micro data presented in Section 2. For this, we assume that the utility cost κ follows a type one extreme value distribution with mean 0 and scale parameter σ_κ equal to $(1/3e^8)$, a small number that implies small dispersion in utility costs. Finally, we choose a small monetary cost equal to $\nu = 0.0012$, equivalent roughly to 0.05 percent of pre-tax wages.

To see how the accuracy of expectations in the model implied by these parameters compares to those observed in SCE data presented in Section 2, we concentrate on expectations of unemployment, a key determinant of savings decisions in the model, for which the SCE elicits expectations in the form of the "percent chance that 12 months from now the unemployment rate in the U.S. will be higher than it is now". The moments that inform our choice of information cost parameters thus relate to this probability of a higher or lower unemployment rate ur 4 quarters in the future. We choose a natural measure of expectational "errors" ε_u for all individuals in the model as the absolute difference between an individual's perceived probability (conditional on her current posterior ϕ_2), denoted $prob(u_{t+4} > u_t|\phi_2)$ and the true probability $prob(u_{t+4} > u_t|z_t)$, conditional on full information about current conditions, as a percentage of the latter

$$\varepsilon_u = \frac{|prob(u_{t+4} > u_t|\phi_2) - prob(u_{t+4} > u_t|z_t)|}{prob(u_{t+4} > u_t|z_t)} * 100.$$

In order to construct an equivalent measure in the data we require a measure of the "true" probability of rising unemployment. For this, we use the consensus estimate from professional forecasts, reflecting ample evidence that professional forecasters provide more accurate predic-

Table III: Expectational errors

	Mean Error	Standard deviation
Model	122	60
SCE Data	123	87

The table shows the mean and standard deviation of expectational errors ε_u in unemployment expectations. Individual errors are calculated as absolute difference between the individual perceived probability that the unemployment is higher in period $t+4$ than in period t and the true probability, as a percentage fraction of the latter. For SCE data, we compare the survey responses to the consensus probability implied by individual density forecasts for the unemployment rate according to the Survey of Professional Forecasters.

tions than those from standard, statistical and economic models (Stark (2010)). In particular, our expectational error in the data equals the absolute difference between the probability of rising unemployment reported by SCE respondents and the consensus probability of a rising unemployment rate from the Survey of Professional Forecasters (an average constructed from individual forecasters' reported forecast densities). Table III compares the mean and standard deviation of the resulting relative errors in the model and the data. The dispersion in expectations is somewhat smaller than in the data, but overall the model replicates the first two moments of expectational errors in the data well.

We compare the results from this benchmark parameterization, with limited information acquisition, to a "full-information" version where all households acquire information every period.[11].

5.3 Benchmark Results

This section presents the main quantitative results of this paper: our benchmark model of dynamic heterogeneous information choice captures key features of the cross-sectional distribution of expectations in US micro data. When we compare its predictions to those from a version with full information, we find substantially increased inequality and aggregate volatility relative to a standard full-information version of the economy.

5.3.1 Information and savings choices

Figure 4 shows the difference in savings rates between informed agents and "uninformed" agents with a 50/50 prior about current productivity, at three different values of the prior for

[11] In practice, we give households a large utility benefit of information acquisition (a negative cost $\kappa < 0$) to ensure universal information acquisition

mean capital. The difference is strongly heterogeneous across the distribution of cash on hand (along the bottom axis), and between unemployed (left column) and employed households (right column). Unemployed households with (close to) zero assets never save, independently of their information. Savings rates of informed households differ strongly from those of the uninformed at low but positive levels of wealth, where permanent income is dominated by future wages and employment prospects, for which the current productivity level is a good predictor. For a given level of current cash on hand, knowing that current aggregate productivity is high (bottom row in Figure 4) thus decreases savings relative to those implied by a diffuse prior, as future wages and employment probabilities are expected to be high, while the reverse holds for low productivity (top row). This effect is stronger for the unemployed (left column of panels), for whom the strong dependence of job-finding rates on aggregate productivity increases the difference in permanent incomes expected by informed and uninformed agents. As cash on hand increases, the difference in savings rates decreases strongly, because savings behavior of the wealthy is less affected by information about wages and employment rates and differences in expected returns only have small effects on savings rates.

The decision rules for information acquisition reflect these differences in savings rates, together with the effects of the monetary cost ν and the stochastic utility cost κ. Decisions are most easily described by the probabilities of information acquisition before the i.i.d. shock κ is realized, depicted in the top row of Figure 5. As expected, information acquisition probabilities are higher at less informative priors. In line with Figure 4, unemployed households (in the top left panel) never acquire information at capital holdings close to zero, where savings are little affected by information and the monetary cost ν is expensive in utility terms. Unemployed agents with low but positive asset holdings and non-neglibible prior uncertainty about current productivity, in contrast, always acquire information, because it substantially improves their savings choices. As cash on hand rises, information acquisition probabilities reach a minimum and then rise again, in line with the utility losses associated with once-and-for-all information choices in Figures 8 and 9 in the previous section.

In contrast to the losses associated with once-and-for-all choices, which differed only little across employment states, the dynamic information acquisition choices of the employed in the top right panel of Figure 5 are markedly different to those of the unemployed. In particular, employed households with little cash on hand, whose savings are less affected by information

Figure 4: Difference in savings policies, informed vs. uninformed

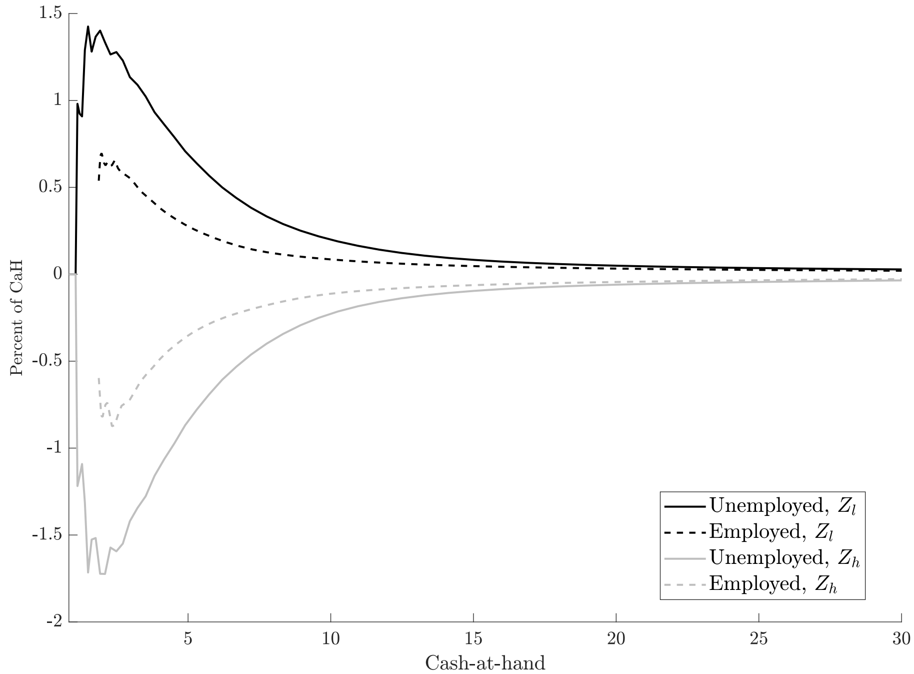

For the unemployed (left column) and the employed (right column), the figure shows the difference between savings by informed agents and uninformed agents ($\phi_1 = 0.5$) at different values of the prior of mean capital $\hat{k}$ (see legend) and different values of individual cash on hand (along the bottom axis), in the benchmark parameterization of the model.

than those of their unemployed peers, only acquire information for low realisations of κ, and when their prior about current aggregate capital is high (implying high future wages). As cash on hand rises, the information acquisition probabilities of the employed and unemployed converge.

5.3.2 Heterogeneity in expectations and average degree of information

The heterogeneity in information acquisition probabilities in Figure 5 translates to strong heterogeneity in expectations across the wealth distribution. In particular, in Figure 6, the errors in expectations about aggregate unemployment 12 months in the future, whose mean and standard deviation we used to discipline the choice of information costs, exhibits an inverse U-shaped relationship with wealth that is similar to that found in SCE data (Figure 1), if less pronounced. Despite its simple information structure, the model thus seems to capture a key

Figure 5: Information acquisition policies

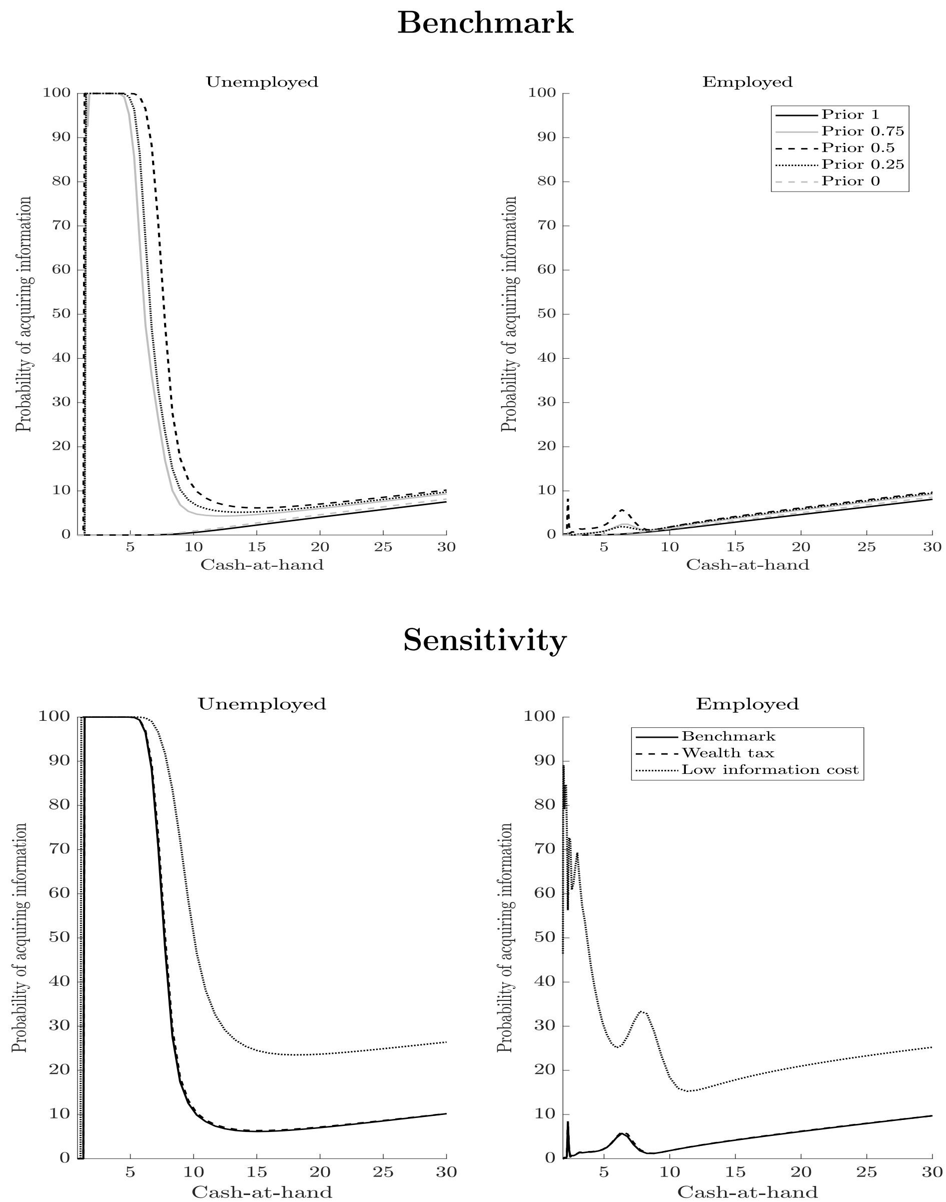

For the unemployed (left column) and the employed (right column), and a mean prior about aggregate capital $\bar{k}$, the figure shows, in the top row, the probability of information acquisition for different values of the prior ϕ_1 and different values of individual cash on hand (along the bottom axis) in the benchmark parameterization of the model. The bottom row depicts information acquisition policies at an uninformed prior about aggregate productivity being high (equal to 0.5), for the benchmark case and alternative parameterisations of the model with low information cost and a wealth tax of 1 percent p.a..

qualitative feature of expectations in US micro data.

Figure 7 shows how the maintained assumption that households' only acquire information about current productivity z_t is not restrictive. In particular, if a household were to acquire

Figure 6: Errors in unemployment expectations

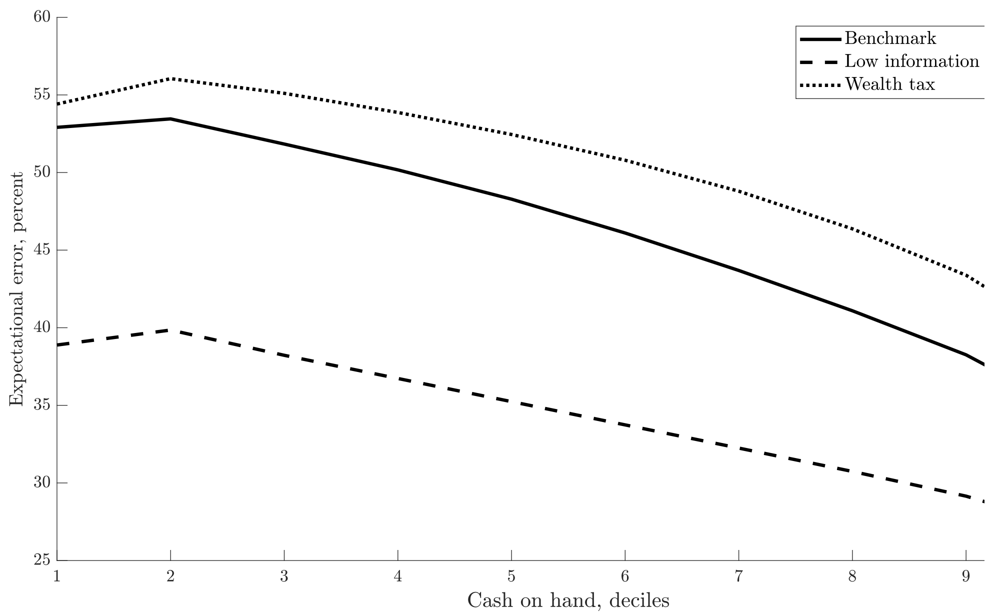

The figure shows errors in unemployment expectations ε_u averaged within deciles of the distribution of individual capital holdings for the benchmark parameterisation (solid line), as well as the economies with low information costs in Section 5.4 (dashed line), and with a 1 percent p.a. wealth tax (dotted line). Individual errors are calculated as the individual perceived probability that the unemployment is higher in period $t+4$ than in period t minus the true conditional probability, as a percentage fraction of the latter.

information about z_t every period, she would make extremely accurate inferences about the current capital stock $\bar{k}_t$ (the gray line in Figure 7), which her prior (the thin black line) follows closely (with a correlation of the two series of about 0.95).[12] On average, however, household priors are slow-moving and less accurate (as indicated by the thick black line in the figure), as only around 15 percent of households decide to acquire information every period. An average agent would thus pay the utility and monetary costs to update her information about the aggregate state of the economy about once every 6 quarters.

[12]In the figure, we start $\hat{k}$ at an arbitrary value of 35, and discard the initial 200 periods to calculate the correlation, to demonstrate that the strong correlation does not depend on an accurate initial point prior.

Figure 7: Mean capital k_t - realization and priors

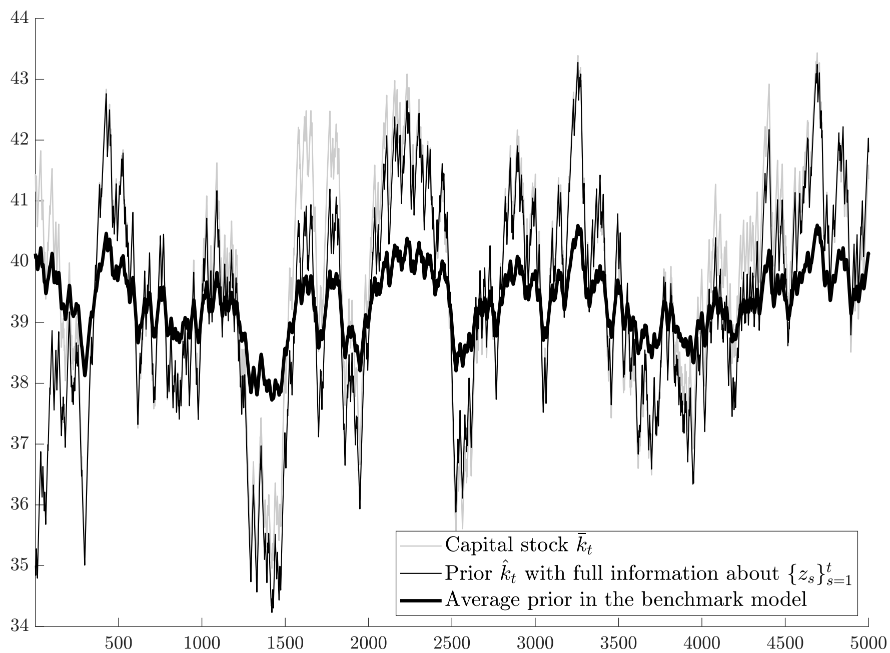

Based on a simulation of the benchmark model, the figure shows time series of the capital stock $\bar{k}_t$ (gray line), the prior about current aggregate capital $\hat{k}_t$ of households who acquire information about the current productivity state every period (thin black line), and the average prior in the cross section (thick black line).

5.3.3 Implications for wealth inequality

Because limited information acquisition changes savings behavior, we would expect the resulting heterogeneity in information to also affect wealth inequality. Table IV summarizes key moments of the wealth distribution, averaged over the ergodic distribution of aggregate states using a long simulation of the economy. While mean wealth is not substantially different with limited information, inequality of wealth holdings is substantially increased. Specifically, the 90/10 (99/1) percentile ratio is about 30 percent higher. The difference in standard deviations is smaller, since they are heavily effected by the extreme right tail of the distribution, where information acquisition probabilities are high in our benchmark model. The reason for this increased wealth inequality is the muted correlation between returns and savings rates with incomplete information. In particular, households' expected labor income rises with current aggregate productivity (which improves future labor market prospects). At the bottom of the wealth distribution, this creates a positive correlation between savings rates and returns as households save more when productivity is high. With limited information about current aggregate conditions, this correlation is less pronounced, reducing wealth holdings at the bottom of the distribution. Households whose wealth is sufficiently high as to make labor income irrelevant, in contrast, save more in times of low productivity. Again, the implied negative correlation between returns and savings rates is less pronounced with imperfect information, increasing the right tail of the wealth distribution.

Table IV: Moments of the wealth distribution, average

	Mean k	St dev	90/10	99/1
Benchmark	38.1	23.2	6.1	18.4
Full information	38.4	20.6	4.7	14.4
Difference, percent	-0.9	12.9	30.0	27.3

The table shows moments of the distribution of individual capital holdings, averaged across a long simulation of the benchmark economy (top row), and its full information counterpart (middle row), plus their percentage difference (bottom row). The moments are the mean (first column), standard deviation (St dev, second column), the ninety-to-ten percentile ratio (90/10, third column), and the ninetynine-to-first percentile ratio (99/1, final column).

5.3.4 Implications for aggregate dynamics

Table V shows how endogenously limited information changes the aggregate dynamics of the economy. Similar to the results in the two-period analysis (and the infinite horizon economy with homogeneous information sets in Appedix 7.5, fluctuations in the capital stock are substantially more pronounced in the benchmark economy, relative to its full-information counterpart. Specifically, the standard deviation of the capital stock is about 40 percent higher, as limited information about its current level dampens the mean-reversion inherent in neoclassical economies, by which lower returns at high capital decrease investment and vice versa. In particular, as Figure 7 shows, the average prior is substantially less volatile than the actual capital stock. With incomplete information, return expectations thus move less, while aggregate output, and thus incomes, move more than under full information. Savings are thus substantially less countercylical than with full information, increasing the volatility of output and other aggregate quantities (although by less than that of capital, as their volatility is dominated by the effect of productivity shocks on current production that are approximately unaffected by information choice).

Table V: Moments of aggregate quantities, average

	St dev K	St dev Y	St dev I	St dev C	Info unemp	Info emp
Benchmark	5.13	3.45	7.77	3.18	0.16	0.14
Full information	3.61	3.16	7.11	2.92	1.00	1.00
Difference, percent	41.92	9.23	9.33	9.02	-84.13	-86.31

The table shows the standard deviations of the time series for natural logarithms of capital (column one), output (column two), investment (column three) and consumption (column four), as well as the average percentage fraction of unemployed (column five) and employed agents (column six) that acquire information in a given period, averaged across a long simulation of the benchmark economy (top row), and its full information counterpart (middle row), all in percent, plus their percentage difference (bottom row).

5.4 Alternative specification of information costs

For our benchmark analysis, we calibrated information costs "top down" to roughly capture the mean and dispersion of expectational errors in US micro data. An alternative would be to estimate the cost of information acquisition at the household level. More generally, the question arises how sensitive our results are to the assumptions about information costs. The fact that incentives to acquire information fall when it becomes cheaper, and more people acquire it should naturally dampen the general equilibrium effect of changes in information costs. While a full sensitivity analysis is beyond the scope of this paper, the bottom row of Figure 5 shows how information acquisition probabilities rise when we set the cost to one half its benchmark level. The average probability of information acquisition rises from about 15 to about 28 percent. This increase is more pronounced at low levels of wealth, so the inverse U-shape that characterises the relationship or expectational errors with wealth is slightly more pronounced than in the benchmark case (see Figure 6). Average errors are also reduced, from 122 to 87 percent.

5.5 Heterogeneous information acquisition and the effects of policies

Through their impact on endogenous information acquisition changes in the economic environment have an additional effect in our benchmark economy relative to the standard case of full-information. In particular, the impact of macroeconomic policies in our model with heterogeneous information acquisition may differ from that in a standard full-information environment. In this section, we provide an example that shows how policies can indeed have an additional, quantitatively important transmission channel by changing information acquisition. For this, we assume that the government, in addition to its labor income tax, also charges a constant wealth tax τ_k on beginning-of-period capital holdings that finances wasteful government spending. Household cash on hand is thus

$$y_t = r_t k_t + (1 - \tau)\bar{l}\epsilon_t w_t + \mu(1 - \epsilon_t)w_t + (1 - \delta - \tau_w)k_t \tag{17}$$

Because the probability of information acquisition is strongly heterogeneous in individual wealth levels (see Figure 5), we would expect the changes in the wealth distribution implied by a positive tax $\tau_w > 0$ to also change average information acquisition. This is exactly what we find: although information acquisition policies at given wealth levels are approximately unaffected by the tax in the bottom row of Figure 5, a wealth tax reduces the mass of households that acquire information every period by reducing average capital holdings (thus moving the distribution of individual capital to the left). Table VI shows that this effect is powerful, amounting to a reduction in the number of individuals who acquire information every period by 15 and 30 percent for a 0.5 and 1 percent p.a. wealth tax, respectively. The fall for the employed is somewhat more pronounced as their information acquisition probability does not show the pronounced hump shape at low asset levels found for the unemployed, and is thus more sensitive to the reduction in wealth implied by the tax. This reduction in information acquisition results in larger expectational errors in Figure 6.

According to Table VI, relative to the full-information economy the reduction in information acquisition does not materially change the fall in average capital implied by the tax (equal to about 5 and 10 percent for the 0.5 and 1 percent tax respectively). Importantly, however, our economy with endogenous information predicts a substantially different impact of the tax on both macroeconomic volatility and inequality. Specifically, Table VI shows that the standard deviations of (the logs of) capital and output are not much affected by the tax with full information. In contrast, by further dampening the mean reversion of the capital stock, the reduction in information acquisition in our benchmark environment *increases* substantially the volatilities of both capital and output, by 8 and 3 percent, respectively, with a 1 percent tax.

Table VI: Impact of a wealth tax

	Mean $\bar{k}_t$	St dev $\bar{k}_t$	St dev Y_t	90/10	99/1	Info acqu.
Limited information, 0.5 %	-5.2	3.6	1.3	-0.5	-1.4	-15.5
Full information, 0.5 %	-5.3	0.7	0.3	-2.3	-2.2	0.0
Limited information, 1 %	-10.0	8.0	2.8	2.0	2.0	-30.4
Full information, 1 %	-10.2	1.3	0.5	-4.3	-7.3	0.0

The table shows percentage changes of equilibrium moments implied by the introduction ofm respectively, a 0.5 and 1 percent p.a. wealth tax ($\tau_k = 0.00125, 0.0025$ for the quarterly tax) in the limited information (rows 1 and 3) and full information (rows 2 and 4) economies. The moments are the mean capital stock (column one), the standard deviations over time of natural logarithms of capital (column two) and output (column three), and 90/10 and 99/1 percentile ratios of the cross-sectional standard deviation of wealth holdings (columns four and five), and the average share of individuals acquiring information every period (column six, for the limited-information economies only).

The difference in the effect of the tax on inequality between the two environments is even more pronounced: by discouraging capital accumulation, the tax reduces wealth inequality strongly with full information. For example, the 90/10 percentile ratio is more than 4 percent lower in the economy with a 1 percent wealth tax. In our benchmark environment, in contrast, there is an offsetting effect because uninformed accumulation policies also dampen mean reversion of individual capital accumulation. With a 0.5 percent tax, the two effects approximately offset each other. With a 1 percent tax, in contrast, the information effect dominates, such that the wealth tax actually *raises* inequality: both percentile ratios increase by 2 percent.

Rather than predictions of the effects of a wealth tax, we view these results as an illustration of how macroeconomic policies may have an additional transmission channel in an environment with endogenous information acquisition, by changing the value of information, and the distribution of agents across heterogeneous information choices. Moreover, these additional effects may be quantitatively important both from a positive and a normative point of view.

6 Conclusion

In this paper we have studied information choice in standard neoclassical economies with heterogeneous agents. The results show how heterogeneity in wealth holdings and employment status naturally implies heterogeneity in the incentives to acquire information. In equilibrium even small information costs therefore imply substantial differences in information acquisition both in the cross-section of households, but also over time as households move through the individual state space. Moreover, individual incentives to acquire information depend on the average degree of information in the economy. This is because informed households save more in times when the aggregate capital stock is low and returns are high, thus dampening business cycles. We show how this implies that rational expectations equilibria are typically not robust to small costs of information acquisition.

These two features, heterogeneous and endogenous incentives for information acquisition, motivate our benchmark analysis of dynamic general equilibrium information choice in an otherwise standard neoclassical economy with idiosyncratic shocks and incomplete markets. With small costs of information, the economy features large heterogeneity in information acquisition, with a substantial fraction of agents optimally not acquiring information in any given period. Importantly, the resulting expectational errors replicate the inverse-U shaped relationship between errors and individual wealth observed in US micro data. Relative to the full-information version of the economy, aggregate dynamics are amplified and wealth and consumption heterogeneity is substantially increased. Moreover, policies have additional effects by changing the average degree of information in the economy, as illustrated by a wealth tax that substantially reduces the mass of households who acquire information every period.

While we believe that the choice of our benchmark environment is sensible, as a one-step deviation from a workhorse model in modern macroeconomics, it is not necessarily without loss of generality. For example, since information can be used freely once acquired, adding additional decisions (such as labor supply, or portfolio choice), may affect the results whenever the resulting optimal behavior varies strongly with the state of the economy. Future research should therefore look at information choice in more general environments. Other dimensions along which our analysis should be made more general are the structure of priors (in particular the point prior about capital), and the dichotomous nature of information choice (that does not allow agents to pay more for better signals, for example).

Our analysis is positive in nature, but raises interesting normative questions. Particularly, information acquisition has obvious externalities in our environment through the implied change in the dynamic properties of prices and aggregate quantities. Does this mean policymakers should subsidize information? Should such subsidies target a particular part of the population? We are interested in these questions, but leave them for future research.

References

Angeletos, G.-M. and C. Lian (2018): "Forward guidance without common knowledge," *American Economic Review*, 108, 2477–2512.

Branch, W. A. and B. McGough (2018): "Heterogeneous expectations and micro-foundations in macroeconomics," in *Handbook of computational economics*, Elsevier, vol. 4, 3–62.

Brock, W. A. and C. H. Hommes (1997): "A Rational Route to Randomness," *Econometrica*, 65, 1059–1095.

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.

Grossman, S. J. and J. E. Stiglitz (1980): "On the impossibility of informationally efficient markets," *The American economic review*, 70, 393–408.

Hellwig, C. and L. Veldkamp (2009a): "Knowing what others know: Coordination motives in information acquisition," *The Review of Economic Studies*, 76, 223–251.

——— (2009b): "Knowing What Others Know: Coordination Motives in Information Acquisition," *Review of Economic Studies*, 76, 223–251.

Krueger, D., K. Mitman, and F. Perri (2016): "Macroeconomics and household heterogeneity," in *Handbook of Macroeconomics*, Elsevier, vol. 2, 843–921.

Krusell, P. and A. Smith (1998): "Income and Wealth Heterogeneity in the Macroeconomy." *Journal of Political Economy*, 106.

Maćkowiak, B. and M. Wiederholt (2009): "Optimal sticky prices under rational inattention," *American Economic Review*, 99, 769–803.

——— (2015): "Business cycle dynamics under rational inattention," *The Review of Economic*

Studies, 82, 1502–1532.

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

McFadden, D. (1974): "The measurement of urban travel demand," *Journal of public economics*, 3, 303–328.

Sims, C. A. (2003): "Implications of rational inattention," *Journal of monetary Economics*, 50, 665–690.

Stark, T. (2010): "Realistic evaluation of real-time forecasts in the Survey of Professional Forecasters," *Federal Reserve Bank of Philadelphia Research Raport, Special Report*, 1.

Townsend, R. M. (1983): "Forecasting the forecasts of others," *Journal of Political Economy*, 91, 546–588.

7 Appendix

7.1 SCE data and sample selection

The SCE is a monthly internet survey that of about 1300 "household heads", defined as the person in a household who owns, is buying or rents the home. Subjects are chosen from the respondents to the Consumer Confidence Survey (CCS), itself based on the universe of US postal adresses, to match demographic targets from the American Community Survey, and remain in the survey for up to 12 months. The SCE core module contains monthly information about households' expectations about key macroeconomic and individual variables, including consumer price and house price inflation, unemployment, interest rates, etc. Importantly, a yearly module also asks the survey respondents for key financial variables, including their net wealth.

We use three expectational variables from the SCE. The first two ask respondents for their best guess for a variable's outcome, in addition to the probability of it falling into a number of bins. The questions are: First, about consumer price inflation: "What do you expect the rate of inflation to be over the next 12 months? Please give your best guess.", and "In your view, what would you say is the percent chance that, over the next 12 months the rate of inflation will be... ", where the word "inflation" is replaced by "the rate of deflation (the opposite of inflation)" for bins that correspond to declining prices. Second, regarding house price inflation: "By about what percent do you expect the average home price to [increase/decrease]? Please give your best guess.", and "And in your view, what would you say is the percent chance that, over the next 12 months, the average home price nationwide will..." We also look at a third question that asks simply for the probability of rising unemployment, "What do you think is the percent chance that 12 months from now the unemployment rate in the U.S. will be higher than it is now?".

For consumer and house prices, we calculate forecast errors by taking the root of the mean squared difference between individual best guesses and the actual (12-month-ahead) outcomes of US consumer price index inflation and inflation of the SP/Case-Shiller 20-City Composite Home Price Index, respectively. We use the measures of interquartile ranges of individual forecasts provided by the SCE. An equivalent procedure is unavailable for individual unemployment forecasts, which only indicate the probability of unemployment rising, during a time period where the US unemployment rate fell monotonically from the high levels that followed the Great Recession. In order to construct errors of individual unemployment forecasts, we compare individual probabilities of rising unemployment to those inferred from the probablistic answers of professional forecasters (in the US Survey of Professional Forecasters, SPF). Specifically, we calculate the mean probability of rising unemployment as the average of the corresponding probabilities of individual professional forecasters (from the probabilistic answers in the variable PRUNEMP). We then subtract this number from the individual probabilities, and scale it by the mean probability of rising unemployment in the sample according to the SPF (to make it comparable to the model-implied probabilities that are calibrated to a different time period).

In addition to these expectational variables we use the following household characteristics: gender, age, dummies that take values 1 if the household head reports to have a college degree, or to participate in the labor market (in the sense that she / he is either employed or unemployed), respectively. We also use a measure of household net financial wealth, which we construct as the difference between a household's total financial assets and non-mortgage debt.[13] We deflate the resulting quantities by the level of the US consumer price index.

We do not perform any sample selection other than dropping households whose median inflation expectations lie in the extreme bins (higher than 12 or lower than -12 percent) respectively.

7.2 Evidence from SCE data

Table VII reports the results for panel regressions that correlate measures of forecast accuracy and precision in SCE data with individual characteristics.

[13]The question about financial assets is "Approximately what is the total current value of your [and your spouse's/partner's] savings and investments (such as checking and savings accounts, CDs, stocks, bonds, mutual funds, Treasury bonds), excluding those in retirement accounts?", that for mortgage debt "Approximately, what is the total amount of outstanding loans against your home(s), including all mortgages and home equity loans?", that for total debt "Approximately, what is the total amount of your [and your spouse's/partner's] current outstanding debt?".

Table VII: Random effects estimates for individual SCE forecasts

	(1) UE, Abs error	(2) Inflation, Abs Error	(3) Inflation, IQR	(4) HP Inflation, Abs Error	(5) HP Inflation, IQR
Male	0.0593	-0.399***	-0.782***	-0.552***	-0.375*
	(1.31)	(-2.89)	(-3.50)	(-3.26)	(-1.75)
College Degree	-0.112**	-0.552***	-0.762***	-0.652***	-0.595**
	(-2.17)	(-3.40)	(-2.90)	(-3.35)	(-2.26)
Participation	0.0356	-0.0626	-0.391*	0.0539	-0.182
	(0.92)	(-0.56)	(-1.83)	(0.27)	(-0.88)
Constant	1.159*	2.124	9.837***	2.590	12.53***
	(1.93)	(1.27)	(2.90)	(1.26)	(3.92)
r2_o	0.0367	0.0910	0.0732	0.0292	0.0350
N	8582	8618	8618	7824	7824

t statistics in parentheses
* $p < .1$, ** $p < .05$, *** $p < .01$

The table shows estimates of a random-effects model for forecast errors and for interquartile ranges (IQR) of the individual forecast distributions in SCE data. Column 1 shows estimates for the errors in individual unemployment forecasts (elicited as the "percent chance that 12 months from now the unemployment rate in the U.S. will be higher than it is now") relative to the equivalent consensus forecast from professional forecasters (see Appendix 2 for detail). Columns 2 and 3 present results for 12 month ahead consumer price inflation, columns 4 and 5 for house price inflation, where forecast errors equal the root of the mean squared difference between individual forecasts and outcomes. All regressions also include a cubic function in age, year-month-dummies, and indicator variables for quintiles of the wealth distribution, whose coefficients are depicted in Figure 1. T-statistics are shown in parentheses. Robust standard errors are used; stars denote conventional significance levels: * ($p < .1$), ** ($p < .05$), *** ($p < .01$).

7.3 Proofs

7.3.1 Proof of Proposition 1

i. $y \leq \underline{y}$ implies that $U'(y) > max_{R,w}\left[RU'(w)\right] \geq E_{R,w}[RU'(w)]$. 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.

ii. The proof follows from the strict concavity of the utility function and the non-degenerate nature of the distribution Φ_1, which implies that unconditional savings choices have strictly positive cost. To see this more formally, choose some $\widetilde{k} > 0$ and consider $\widetilde{y} = \widetilde{k} + \frac{1}{E[\frac{\beta R}{w+R\widetilde{k}}]}$ such that $\widetilde{k}$ solves the Euler equation without information acquisition for an individual with period one income $y = \widetilde{y}$. Note that in any state of nature $\{R, w\} \in \Psi$ the utility implied by informed savings choices is never smaller than that implied by the uninformed savings choice $\widetilde{k}$. To see how it is strictly greater in some states with positive probability, note that the partial derivatives of k with respect to w and R implied by equation (13) are strictly non-zero when k is interior. Since the distribution $\Phi_1(Y, R)$ is not degenerate, there exist thus $\{\widehat{w}, \widehat{R}\} \in \Psi$ with $\Phi_1(\{\widehat{w}, \widehat{R}\}) > 0$, such that the optimal savings choice $k(\widehat{w}, \widehat{R}) - \widetilde{k} = \epsilon$ with $abs(\epsilon) > 0$ when the individual receives information $\{\widehat{w}, \widehat{R}\}$. The strict concavity of period utility implies that perturbing the optimal saving by a non-zero amount ϵ strictly reduces utility in state $\{\widehat{w}, \widehat{R}\}$. There is thus a κ small enough such that the individual optimally acquires information.

iii.a The strategy of the proof is to bound the loss of not acquiring information $\mathbb{L}(y) = U_{noninf}(y) - U_{INF}(y)$ from above. For this, we first compute an upper bound for the difference in savings choices consistent with (13) and (14), 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 $\widehat{b}$. 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 bounded 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:

$$
\begin{aligned}
\mathbb{L} \geq \;& log\left(\frac{1}{1+\beta}(y+\frac{\bar{w}}{\underline{R}})\right) + \beta log\left(\frac{\bar{R}\beta}{1+\beta}(y-\frac{\bar{w}}{\beta\underline{R}})\right) && (18)\\
& -log\left(\frac{1}{1+\beta}(y)\right) - \beta log\left(\frac{\bar{R}\beta}{1+\beta}(y)\right) && (19)\\
& +log\left(\frac{1}{1+\beta}(y)\right) + \beta log\left(\frac{\underline{R}\beta}{1+\beta}y+\bar{w}\right) && (20)\\
& -log\left(\frac{1}{1+\beta}(y+\frac{\bar{w}}{\underline{R}})\right) - \beta log\left(\frac{\underline{R}\beta}{1+\beta}(y+\frac{\bar{w}}{\underline{R}})\right) && (21)\\
= \;& log\left(1+\frac{\bar{w}}{\underline{R}y}\right) + \beta log\left(1-\frac{\bar{w}}{\beta\underline{R}y}\right) + log\left(1-\frac{\frac{\bar{w}}{\underline{R}y}}{1+\frac{\bar{w}}{\underline{R}y}}\right) + \beta log\left(1+\frac{\frac{\bar{w}}{\underline{R}y}}{\beta(1+\frac{\bar{w}}{\underline{R}y})}\right) && (22)\\
< 0 \;& && (23)
\end{aligned}
$$

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 $\widehat{y}$ such that $\kappa > -\mathbb{L}$ for any $y > \widehat{y}$.

iii.b Note that for every R, w pair, utility is strictly concave in k, so there is a unique $k^\star$ that maximises utility with information acquisition. Expected utility across the distribution of R, w is also strictly concave in k, with a unique maximand $k^{\star\star}$. Moreover, with $\gamma \neq 1$, $k^{\star\star} \neq k^\star$ for at least one R, w pair. With $\gamma \neq 1$, the associated savings rates $\hat{k}^{\star\star} = \frac{k^\star}{y}, \hat{k}^{\star\star} = \frac{k^{\star\star}}{y}$ converge to different constants as y grows large. Moreover, for every R, w pair, any $\hat{k}$ not equal to the optimal $\hat{k}^\star$ yields strictly lower utility. So there is a strict utility loss from uninformed savings choice $k^{\star\star}$ that approaches a constant fraction of $\frac{y^{1-\gamma}}{1-\gamma}$ as y grows large. This yields the result.

7.3.2 Proof of Proposition 3:

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 (13), (3) and (4) that $\frac{\delta k(y,K)}{\delta K} < 0$. The linearity of the implied savings rule in (13), 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 (2009b) to show the value of acquiring information 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.

7.4 The household problem with extreme value shocks

This Section of the online appendix provides detail about the version of the household problem (6) to (11) we use in our quantitative analysis. There are two differences with respect to the benchmark problem: first, households decide each period whether or not to acquire information about the exogenous value of current productivity z_t, which they use to update their point prior about mean capital $\hat{k}_t$. Second, households experience utility cost shocks associated with both choices, acquiring information and not, equal to, respectively, κ_{INF} and κ_{NOINF}. These shocks follow independent, identical type-1 extreme value distributions with mean 0 and scale parameter σ_κ.

Stage 1

In the first stage of their dynamic program households choose information whether to acquire information or not to acquire information about z

$$V_1(y, \epsilon; p, \hat{k}) = max\left[V_{INF}(y, \epsilon; p, \hat{k}) + \kappa_{V_{INF}}, V_{NOINF}(y, \epsilon; p, \hat{k}) + \kappa_{NOINF}\right] \tag{24}$$

where y is cash on hand, p is the prior probability of $z_t = z_h$, and V_{INF} and V_{NOINF} denote the values associated, respectively, with buying information and not buying informaiton about z.

From the McFadden (1974) formula, the expected value across realisations of the extreme value shocks, conditional on given values of $V_{INF}(y, \epsilon; p, \hat{k})$ and $V_{NOINF}(y, \epsilon; p, \hat{k})$ is

$$E[V_1|V_{INF}, V_{NOINF}] = \frac{\gamma}{\alpha} + \frac{1}{\alpha}\left[log(e^{V_{INF}} + e^{V_{NOINF}})\right] \tag{25}$$

and

$$\mathbb{P}_{INF}(y, \epsilon; p, \hat{k}) = \frac{e^{V_{INF}}}{e^{V_{INF}} + e^{V_{NOINF}}} \tag{26}$$

where $\mathbb{P}_{INF}(y, \epsilon; p, \hat{k})$ denotes the ex-ante probability of information choice.

Note that we can define the expected value of utility shocks κ as

$$
\begin{aligned}
E_\kappa(y,\epsilon;p,\hat{k}) \equiv \mathbb{P}_{INF}(y',\epsilon';p_{z'|z},\hat{k}')E[\kappa_{INF}|V_{INF}+\kappa_{INF} > V_{NOINF}+\kappa_{NOINF}] \\
+(1-\mathbb{P}_{INF}(y',\epsilon';p_{z'|z},\hat{k}'))E[\kappa_{NOINF}|V_{INF}+\kappa_{INF} \leq V_{NOINF}+\kappa_{NOINF}] \\
= \frac{\gamma}{\alpha}+\frac{1}{\alpha}\left[log(e^{V_{INF}}+e^{V_{NOINF}})\right] \\
-\mathbb{P}_{INF}(y',\epsilon';p_{z'|z},\hat{k}')V_{INF}-(1-\mathbb{P}_{INF}(y',\epsilon';p_{z'|z},\hat{k}')V_{NOINF}
\end{aligned}
\tag{27}
$$

where we suppress the dependence of Values on $y,\epsilon;p,\hat{k}$ for brevity.

Stage 2

In the second stage, households choose consumption and savings given their prior or acquired information. Importantly, the second stage problem is not recursive in the state variables $y,\epsilon;p,\hat{k}$. This is because agents form expectations about future events knowing the dependence between cash-on-hand y and aggregate productivity z. To illustrate, suppose that all agents are employed when $z=z_h$ and all are unemployed when when $z=z_l$. Take an agent who has an uninformative prior $p=0.5$ and does not buy information, such that her posterior equals 0.5 also. She knows that if she is employed tomorrow, she is in the high-productivity state ($z=z_h$), and therefore likely to remain at employment and high aggregate productivity. But her prior next period is still $p=0.5$. So only using the prior as a function of previous information about aggregate productivity as a state variable imposes forgetting about the correlation of y and z next period, and thus the wrong probability distribution about *following* periods when forming expectations in stage two.

Instead, to form expectation, agents consider a conditional value $V_{INF,cond}(y,\epsilon;z;\hat{k})$ that includes z as a state variable.

$$
V_{INF}(y,\epsilon;p,\hat{k}) = pV_{INF,cond}(y,\epsilon;z_h;\hat{k})+(1-p)V_{INF,cond}(y,\epsilon;z_l;\hat{k}) \tag{28}
$$

$$
\begin{aligned}
V_{INF,cond}(y,\epsilon;z;\hat{k}) = \; & max_{c,k'}[\frac{c^{1-\gamma}-1}{1-\gamma}+\beta\ \Pi_{z',\epsilon'|z,\epsilon}[\\
& \mathbb{P}_{INF}(y',\epsilon';p_{z'|z},\hat{k}')V_{INF,cond}(y',\epsilon';z_h;\hat{k}') \\
& +(1-\mathbb{P}_{INF}(y',1;p_{z'|z},\hat{k}'))V_{NOINF,cond}(y',\epsilon';z';\Pi^z_{z'|z},\hat{k}') \\
& +E_\kappa(y',\epsilon';\Pi^z_{z'|z},\hat{k}')]^z
\end{aligned}
\tag{29}
$$

$$V_{NOINF}(y,\epsilon;p,\hat{k}) = max_{c,k'}\{pV_{NOINF,cond}(y,\epsilon;z_h;p,\hat{k}) + (1-p)V_{NOINF,cond}(y,\epsilon;z_l;p,\hat{k})\} \quad (30)$$

$$\begin{aligned} V_{NOINF,cond}(y,\epsilon;z;p,\hat{k}) &= \frac{c^{1-\gamma}-1}{1-\gamma} + \beta\Pi_{z',\epsilon'|z,\epsilon}[\\ &\quad \mathbb{P}_{INF}(y',\epsilon';p',\hat{k}')V_{INF,cond}(y',\epsilon';z';\hat{k}') \\ &\quad +(1-\mathbb{P}_{INF}(y',\epsilon';p',\hat{k}'))V_{NOINF,cond}(y',\epsilon';z';p',\hat{k}') \\ &\quad +E_\kappa(y',\epsilon';p',\hat{k}')] \end{aligned} \quad (31)$$

$$\begin{aligned} s.t. \quad c+k'+\nu &= y \\ y' &= r'(z',\hat{k}')k + (1-\tau)\bar{l}\epsilon w'(z',\hat{k}') + \mu(1-\epsilon)w'(z',\hat{k}') + (1-\delta)k' && (32) \\ \widehat{k}' &= pG(z_h,\hat{k}) + (1-p)G(z_l,\hat{k}) && (33) \\ p' &= p\Pi^z_{z'_h|z_h} + (1-p)\Pi^z_{z'_h|z_l} && (34) \\ c,k' &\geq 0 && (35) \end{aligned}$$

Note that the McFadden (1974) formula does not allow us to condition the expected value $E[V_1]$ prior to information acquisition on anything but the state variables that are known at the time of decision. In (29) and (31) we have therefore split $E[V_1|V_{INF},V_{NOINF}]$ into the probability-weighted conditional values and the conditional expectation of utility shocks $E_\kappa(y,\epsilon;p,\hat{k})$.

7.5 Once-and-for-all information choice and the KS equilibrium

In this section we study KS-type equilibria where all households *every period* acquire the same information about z_t and $\bar{k}_t$, implying that Ω_{it} is constant across households and time, equal to $\widetilde{\Omega}$. Conditional on this "equilibrium" information choice, we calculate the expected utility implied by "deviations", where a household *every period* makes a different information choice Ω to form expectations about the future. We then calculate the expected utility losses from using only a subset of information relative to the comprehensive information set $\Omega_{max} = \{z,\overline{k}\}$) (which KS find to allow extremely accurate predictions of future prices).[14]

The left panel of Figure 8 considers $\widetilde{\Omega} = \{z,\overline{k}_t\}$ (the KS benchmark where G takes the form of two simple autoregressions of $log(\overline{k'})$ on $log(\overline{k})$ conditional on z). It presents the utility loss for the unemployed of taking savings choices without any knowledge of the current state of the

[14]Appendix 7.6 provides more detail about how we calculate utility losses.

economy (such that $\Omega = \emptyset$ and agents only use average transitions and the unconditional mean of capital in their forecasts for the future). Relative losses are markedly different at different levels of cash on hand, following a pattern similar to that observed in the simple two-period economy (in the right-hand panel of Figure 2). Losses are highest at low but positive asset holdings, where savings choices are unconstrained and future income is dominated by wages that are predicted substantially better with information about current aggregate productivity (whose persistence makes it a good predictor for future separation rates) and current mean capital (which predicts the level of future wages). As wealth rises, the difference between informed and uninformed savings policies (the dashed lines) decreases. Losses therefore also decline, reaching a minimum when wealth is most optimally diversified across human and financial capital given the variance-covariance structure of wages and returns to capital. With high job-finding rates of 50 percent on average, the losses for the unemployed in the figure are similar to those of the employed (but lower at levels of asset holdings close to zero, where the unemployed's current savings choice is constrained and expected losses only arise from uninformed choices in future unconstrained periods).

Figure 8: Savings and utility with high and low information

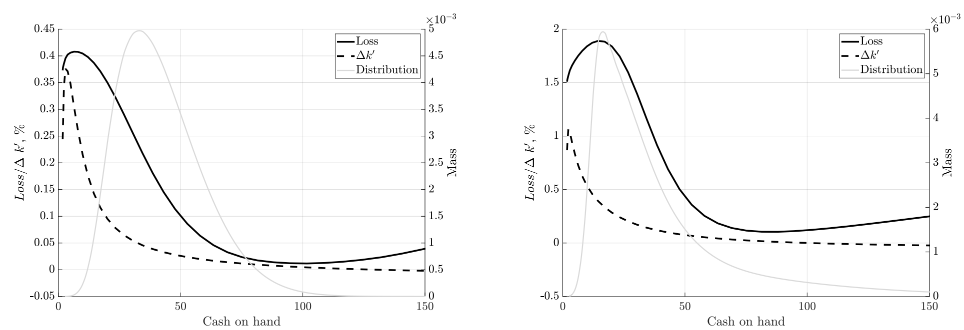

(a) High-information economy ($\widetilde{\Omega} = \{z, \bar{k}\}$) (b) Low-information economy ($\widetilde{\Omega} = \emptyset$)

The Figure shows results for the unemployed in two KS-type economies with constant information sets for all households equal to $\widetilde{\Omega} = \{z, \bar{k}\}$ (the KS benchmark, left-hand panel) and $\widetilde{\Omega} = \emptyset$ (the "low-information" economy, right-hand panel), when aggregate productivity is low ($z = Z_l$) and mean capital equals its median value in a long simulation of the economy. The figure depicts $\Delta k' = (k'(k, \epsilon; \{z, \bar{k}\}) - k'(k, \epsilon; \emptyset))/y * 100$ (dashed lines), the difference in individual savings between using the maximum information about the state of the economy and using no information (implying forecasts of aggregate capital equal its unconditional mean), as a percentage share of cash on hand. It also shows the implied difference in expected utility (solid lines), and the average cross-sectional distribution of households conditional on aggregate productivity averaged within a 2.5-percentile band around the median of mean capital $\bar{k}_t$.

Losses in the left panel of Figure 8 are small on average, and negligible for some households in the distribution. This implies that the KS equilibrium is not robust to small information costs, which would make at least some households deviate from $\widetilde{\Omega}$. Does this mean that all agents should be expected to use less information in equilibrium when there are small costs of information? Not necessarily. The right panel of Figure 8 depicts expected utility losses from uninformed consumption and savings choices equivalent to those in the left panel but in a low-information economy (where $\widetilde{\Omega} = \emptyset$, such that all (other) households only take unconditional decisions). While the relationship of losses with individual wealth maintains its shape, losses are now substantially larger on average. They are also more variable, by up to a factor of two, across aggregate states, as shown in Figure 9 in appendix 7.6. Table VIII confirms this: average losses are between 5 and 6 times higher in the low-information economy (column 3, where $\widetilde{\Omega} = \emptyset$) relative to the KS benchmark ($\widetilde{\Omega} = \{z_t, \bar{k}_t\}$, column 2). Figure 10 in the appendix shows why: the capital stock is substantially more volatile and more persistent when households do not condition their savings choices on current productivity and the current level of capital. In other words, like in the two-period economy of Section 4, uninformed savings choices in equilibrium strongly dampen the mean reversion inherent in neoclassical economies (whereby the higher returns implied by a lower capital stock increase savings in bad times, and vice versa for good times). The implied widening of the capital distribution around its average makes uninformed savings choices more costly. And the increase in persistence makes information about the current level of capital even more valuable to predict the future.

Rows 2 and 3 of Table VIII report losses for two alternative information sets: when households know the current level of productivty ($\Omega = \{z_t\}$, second row) and use it to condition their employment risk and expectations about the future capital stock (effectively replacing the unconditional with the conditional expectation $E[\bar{k}_{t+1}|z_t]$), losses are reduced only slightly. So information about *current* productivity per se is not extremely valuable (as in Proposition 3). When households form expectations about capital using an unconditional autoregression for the mean capital stock (corresponding to $\Omega = \{\bar{k}_t\}$, third row of Table VIII) the expected losses are strongly reduced (and, in relative terms, even slightly smaller in the low-information economy with $\widetilde{\Omega} = \emptyset$). So information about the current level of capital is particularly valuable for savings choices. Finally, maximum losses of not acquiring information in the low-information economy ($\widetilde{\Omega} = \emptyset$, in column 4 of Table VIII), exceed that in the KS benchmark ($\widetilde{\Omega} = \{z_t, \bar{k}_t\}$, column 5) by a similar factor as expected losses.

Table VIII: Utility losses in percent CE

	Expected Losses		**Maximum Losses**	
	$\widetilde{\Omega} = \{z, \bar{k}_t\}$	$\widetilde{\Omega} = \emptyset$	$\widetilde{\Omega} = \{z, \bar{k}_t\}$	$\widetilde{\Omega} = \emptyset$
$\Omega = \emptyset$	0.1764	0.9812	0.7063	3.9455
$\Omega = \{z\}$	0.1538	0.9044	0.5514	3.6597
$\Omega = \{\bar{k}_t\}$	0.0957	0.0672	0.1499	0.1592

The Table presents expected and maximum losses from using different individual information sets Ω (indicated in the first column) rather than a comprehensive information set $\Omega_{max} = \{z, \bar{k}\}$, for different specifications of the information used by all other households $\widetilde{\Omega}$ (indicated in the top row). Expectations and maxima are taken across the 2.5-97.5 percentile range of the ergodic distribution of aggregate and individual states.

Just as in Section 4, these results imply that for a range of utility costs there is no KS-type equilibrium with constant homogeneous information $\widetilde{\Omega}$: Information costs are too high for the KS-benchmark, but too low for a low-information equilibrium.

Result 1 *In the quantitative, infinite-horizon economy with constant homogeneous information choice* $\widetilde{\Omega} \in \{\emptyset, \{z_t, \bar{k}_t\}\}$, *there is a non-empty range of information costs such that there exists no equilibrium.*

In fact, any information cost structure whose permanent consumption equivalent exceeds the minimum utility loss in the KS economy (with $\widetilde{\Omega} = \{z_t, \bar{k}_t\}$) but does not exceed the maximum utility loss in low-information economy ($\widetilde{\Omega} = \emptyset$) has this property.

Note that we considered losses relative to a benchmark "comprehensive" information set equal to $\Omega_{max} = \{z_t, \bar{k}_t\}$ in this section. This is motivated by KS, who show that, when all agents use that information set, considering more information (in the form of additional moments of the cross-sectional distribution) improves the in-sample fit of the predictive regression for the capital stock by only a small amount. They also argue that the welfare gains are "vanishingly small" (p. 878). We confirmed this result in several exercises: The average expected welfare gain from increasing the information set to also contain the variance of individual assets, i.e. $\Omega = \{z_t, \bar{k}_t, var(k)_t\}$, are less than one hundredth of a percent of permanent consumption. In other words, the relevant information choice in standard neoclassical heterogeneous-agent economies is about giving up information relative to KS, not about adding more.

7.6 KS-type equilibria: Further details

This section provides more detail on the KS-type equilibria of Section 7.5. We first describe how we calculate the utility loss from less-informed decisions. We first calculate optimal decision rules for consumption and savings for "deviations" from the equilibrium information set $\widetilde{\Omega}$, where a household *every period* uses a different information set Ω to form expectations about the future only conditional on the information they have. For example, for $\Omega = \bar{k}$, agents would use an unconditional law of motion for mean capital. We then 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}, transformed into units of permanent consumption (and abstracting form information costs). Losses thus equal $CE_{\Omega}(y, \epsilon; \Omega_{max}) = (1 - \beta)\left[(1-\gamma)\left(V_{\Omega}(y, \epsilon; \Omega_{max}) - V_{\Omega_{max}}(y, \epsilon; \Omega_{max})\right) + 1\right]^{\frac{1}{1-\gamma}}$, $\Omega \subset \Omega_{max}$, $V_{\Omega}(k, \epsilon; \Omega_{max})$ equals the discounted utility that 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 we suppress the dependence on the assumed information set $\widetilde{\Omega}$ used by all other households for simplicity.

Figure 9 presents the utility loss for the unemployed of taking savings choices without any knowledge of the current state of the economy (such that $\Omega = \emptyset$ and agents only use average transitions and the unconditional mean of capital in their forecasts for the future), in the no-information equilibrium, where $\widetilde{\Omega} = \emptyset$. Figure 10 presents the time series of mean capital $\bar{k}_t$ in the KS-equilibrium ($\widetilde{\Omega} = \{z_t, \bar{k}_t\}$) and the no-information equilibrium ($\widetilde{\Omega} = \emptyset$).

Figure 9: Utility losses from uninformed savings with homogeneous constant information set $\widetilde{\Omega} = \emptyset$

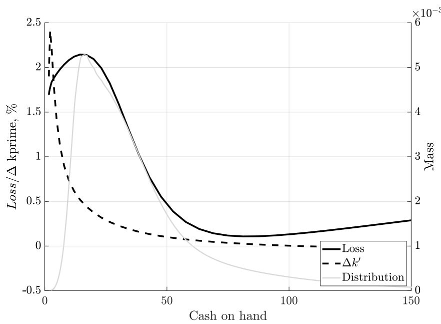

(a) $z = Z_l$, low $\bar{k}$

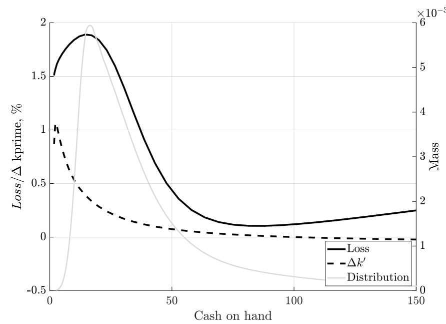

(b) $z = Z_h$, low $\bar{k}$

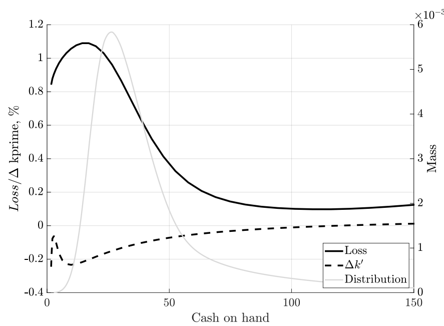

(c) $z = Z_l$, high $\bar{k}$

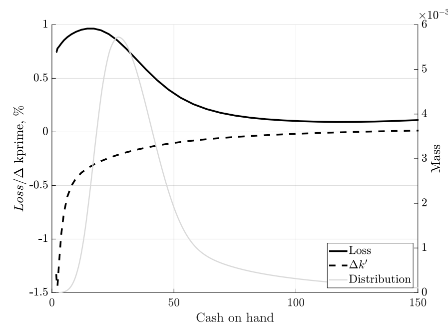

(d) $z = Z_h$, high $\bar{k}$

For the low-information economy ($\widetilde{\Omega} = \emptyset$), the figure presents differences in policy functions $k'(k, \epsilon; \Omega)$ for $\Omega = \emptyset$ (implying forecasts of aggregate capital equal its unconditional mean) and $\Omega = \{z, \bar{k}\}$, respectively (dashed lines), utility losses $CE_{\Omega}(k, \epsilon; \Omega_{max})$ for $\Omega_{max} = \{z, \bar{k}\}$ and $\Omega = \emptyset$ (solid lines), and the average cross-sectional distribution of households conditional on aggregate productivity and averaged within a 2.5-percentile band around high and low $\bar{k}$, corresponding to the first and third quartile of the distribution of mean capital, respectively. The figure concentrates on the employed ($\epsilon = 1$).

Figure 10: Aggregate time series of capital in KS-type economies with homogeneous constant information set $\widetilde{\Omega}$

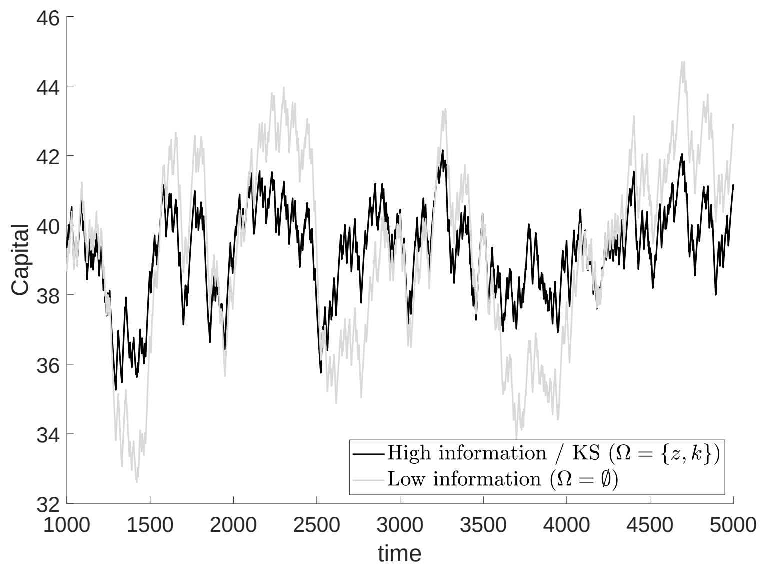

The figure presents time series of the aggregate capital stock from a simulation of the full-information (KS) economy ($\widetilde{\Omega} = \{z_t, \bar{k}_t\}$, black line) and the information with uninformed savings ($\widetilde{\Omega} = \emptyset$, gray line).