Household Consumption and Dispersed Information*

Jonathan J. Adams[†]
University of Florida

Eugenio Rojas[‡]
University of Florida

October 14, 2022

WORKING PAPER
Link to Most Current Version

Abstract

We study the effects of aggregate income shocks in a small open economy heterogeneous agent model. By introducing a standard information friction, we are able to explain two patterns of small economies experiencing large income changes: (1) excess volatility in consumption and (2) household consumption elasticities that have low correlation with income. With a standard dispersed information structure, households cannot distinguish aggregate income shocks from idiosyncratic ones. Therefore their consumption responds excessively to aggregate income changes, which they forecast as likely to be more persistent than they would if they had full information. We demonstrate that this effect occurs at all points in the income distribution, lowering the correlation of the consumption elasticity with income. Finally, we corroborate our central mechanism using survey data on household expectations of their future income.

JEL-Codes: D84, E21, E32
Keywords: Heterogeneous Agents, Incomplete Information, Heterogeneous Beliefs, Business Cycles, Consumption Volatility

*We thank Robert Ainsworth, German Bet, Yu-Ting Chiang, Elias Dinopoulos, Joel Flynn, Rafael Guntin, Alexandre Kohlhas, Felipe Saffie, David Sappington, and seminar participants at the University of Florida and SED 2022 for helpful comments and suggestions. All errors are our own.

[†]Department of Economics, University of Florida, 333 Matherly Hall, PO Box 117140, Gainesville, FL 32611. Email: adamsjonathan@ufl.edu.

[‡]Department of Economics, University of Florida, 331 Matherly Hall, PO Box 117140, Gainesville, FL 32611. Email: erojasbarros@ufl.edu.

1 Introduction

Why is aggregate consumption so volatile? This is a classic question in macroeconomics. Traditional models with full information and rational expectations (FIRE) predict that consumption will be relatively inelastic to aggregate income changes, as households smooth consumption (Kydland and Prescott, 1982). However, consumption volatility is much higher than expected, even more so in developing open economies (Aguiar and Gopinath, 2007). Possible explanations abound, including financial frictions: when agents face collateral or borrowing constraints, they become more elastic to transitory income shocks (Carroll, 2001). However, recent evidence suggests that standard household financial frictions cannot be the main mechanism. The financial friction mechanism gives a clear prediction: the consumption elasticity with respect to an aggregate income shock should decrease with household income. Yet, evidence for this pattern is weak; Guntin et al. (2020) (hereafter GOP) document that the consumption elasticity to large macroeconomic shocks remains high across the entire income distribution, using household-level consumption data from five small open economies.[1] They conclude that consumers of all income levels appear to respond as if changes to aggregate income are permanent shocks.

We propose an explanation that is consistent with the finding by GOP: households face an information friction. Specifically, households are unable to accurately distinguish idiosyncratic income shocks from aggregate income shocks. Crucially, the idiosyncratic component of household income is more persistent than aggregate income. Therefore, if household income rises due to an aggregate shock, households will expect their income improvement to be more persistent than if they had full information. Households respond by increasing consumption by more than if they correctly predicted that the income improvement would be short-lived.

We study this mechanism in a tractable heterogeneous agent model with dispersed information. Households in a small open economy receive stochastic income, and solve a standard consumption-savings problem to mitigate their income risk and smooth consumption over time. Income is determined by two stochastic components: an idiosyncratic income process, and a less persistent aggregate income process. Households have rational expectations,

[1] Like Guntin et al. (2020), we are careful to distinguish the distribution of consumption elasticities to aggregate income shocks from the distribution of marginal propensities to consume (MPCs), which are elasticities to unexpected idiosyncratic windfalls. Empirical evidence is clear that the MPCs to transitory income shocks are decreasing with income and especially household liquidity. For example, Johnson et al. (2006) document this pattern in the response to 2001 tax rebates, Parker et al. (2013) do the same for 2008 stimulus checks, as do Chetty et al. (2020) for COVID-19 relief. These relationships with income are consistent with standard theories, including our own. Kaplan and Violante (2014) augment a standard theory to explain why liquidity matters more than income or wealth.

but incomplete information. They do not directly observe aggregate income shocks, and so cannot accurately distinguish between the two components, which distorts their consumption and savings decisions. We find that the volatility of aggregate consumption growth is one third larger when agents face the information friction. The additional risk induces a stronger precautionary savings motive, so average savings is 43% higher than the full information baseline. And while consumption is more volatile, the elevated wealth also increases the average consumption level.

In the cross-section, we also recover results that are in line with GOP's finding of relatively homogeneous consumption responses to income shocks. Specifically, we document that the consumption elasticity to aggregate income is larger and more homogeneous across the income distribution when agents face information frictions, where the slope of the consumption elasticity across log income levels is nearly zero. This is the first of several results that require studying incomplete information and heterogeneous agents in a unified framework. Moreover, we find the the information friction and financial friction interact in a variety of rich ways: the frictions jointly attenuate inequality dynamics, reduce the sensitivity of the wealth distribution to the borrowing constraint, reverse the relationship between idiosyncratic risk and the aggregate consumption elasticity, and generate endogenous correlations between aggregate forecasts and wealth.

Is our information friction realistic? To answer this question, we employ survey data on household expectations, and document evidence corroborating our central mechanism. The main implication of the information friction is that the response of households forecasts of their own income should respond with the same elasticity to aggregate and idiosyncratic income shocks. If instead households have full information, then their forecasts will be *less elastic* to aggregate than to idiosyncratic shocks. This is a testable prediction. We employ data on household forecasts from the NY Fed's Survey of Consumer Expectations, and decompose household income into aggregate and idiosyncratic components at the state level. Our results are clear: household forecasts are *at least as elastic* to aggregate shocks as to idiosyncratic shocks. Thus we confirm that GOP's characterization of consumption behavior applies to household expectations as well: consumers' forecasts respond to aggregate income changes as if they are more persistent than they really are.

Our paper contributes to several strands of the literature. First, we join a small but promising new literature that synthesizes incomplete information theories with heterogeneous agent models. Broer et al. (2021) and Broer et al. (2022) depart from the standard FIRE structure by introducing a rational inattention decision that is endogenously heterogeneous.[2]

[2]Other papers depart from FIRE in heterogeneous agent models by relaxing rational expectations rather than of full information. For example, Auclert et al. (2020) and Carroll et al. (2020) assume agents have

Angeletos and Huo (2021) show that myopic effects of information frictions are exacerbated by the MPC distribution typical of HANK models. Gallegos (2022) extends Bilbiie (2020) to study a linear HANK model with dispersed information.[3]

Second, we contribute to the empirical literature that examines the effects of aggregate shocks on forecasts of idiosyncratic variables. Many papers study forecasts of aggregate quantities, about which survey evidence is bountiful. But less works exists for households' or firms' forecasts of factors that are specific to them. Among this group, Andrade et al. (2022) study how firms' forecasts of their own prices and production respond to aggregate and industry-level shocks, finding support for the standard dispersed information structure. Adams-Prassl et al. (2022) study workers' perceptions of the returns to searching for a job, and they document a channel through which incomplete information affects worker behavior: workers who are more optimistic about the macroeconomy believe they are more likely to receive a job offer.

Third, we contribute to a class of heterogeneous agent models attempting to understand large consumption responses to income. While we follow GOP and study the consumption-income elasticity to aggregate shocks, many more papers focus on explaining large MPCs. We consider our explanation of the consumption-income elasticity to be complementary to this literature, which Kaplan and Violante (2022) survey. While our simple model cannot explain the cross-sectional evidence on consumption out of idiosyncratic windfalls, our large consumption elasticities to aggregate shocks are relevant for many of the same aggregate applications.[4] One advantage of our approach is that we attain large consumption-income elasticities, while maintaining a realistic wealth distribution that does not suffer from the "missing middle" problem that Kaplan and Violante identify as plaguing most single asset models that otherwise achieve large MPCs.

The remainder of the paper is organized as follows. In Section 2 we describe our model, including the structure and intuition for the information friction. Section 3 describes our main results and equilibrium behavior in the model. In Section 4 we explore the interactions between the information and asset market frictions. Section 5 documents the empirical evidence corroborating our information friction. Section 6 concludes.

---

sticky expectations, in the style of Mankiw and Reis (2002) and Carroll (2003).

[3] Angeletos and Lian (2016) survey the broader literature of incomplete information in macroeconomics. See Heathcote et al. (2009), Quadrini and Ríos-Rull (2015), Krueger et al. (2016) for broad surveys of the household heterogeneity models, and Kaplan and Violante (2018) for a more recent survey that includes HANK features.

[4] This includes monetary policy (Kaplan et al., 2018), fiscal policy (Auclert et al., 2018), and aggregate shocks in general (Bilbiie, 2020) among others.

2 Model

In this section we describe our baseline model. Heterogeneous agents trade risk-free assets in a small open economy in order to self-insure against income risk and smooth consumption. The agents face a standard friction: dispersed information in the style of Lucas (1972) that prevents them from observing the aggregate state of the economy.

2.1 Households

There is a unit measure of identical and infinitely lived households. Households are indexed by i and time is indexed by t.

The household's preferences over current and future consumption are represented by the utility function

$$E_{i,t}\left[\sum_{s=0}^{\infty}\beta^s\frac{C_{i,t+s}^{1-\gamma}-1}{1-\gamma}\right] \tag{1}$$

where $C_{i,t}$ is the household's consumption in period t, β is its discount factor, and γ is the coefficient of relative risk aversion. The expectation operator $E_{i,t}$ is conditional on household i's information set $\Omega_{i,t}$.

The household receives stochastic income $Y_{i,t}$, which in logs is the sum of a mean zero idiosyncratic component $Y_{i,t}^I$ and a common aggregate component Y_t^G:

$$\ln Y_{i,t} = \ln Y_{i,t}^I + \ln Y_t^G \tag{2}$$

The idiosyncratic and aggregate components each follow an AR(1) process:

$$\ln Y_{i,t}^I = \rho_I \ln Y_{i,t-1}^I + u_{i,t}^I \qquad \ln Y_t^G = \rho_G \ln Y_{t-1}^G + u_t^G \tag{3}$$

with $u_{i,t}^I \sim N(0,\sigma_I^2)$, $u_t^G \sim N(0,\sigma_G^2)$, $\rho_I \in (0,1)$, and $\rho_G \in (0,1)$. Crucially, we assume $\rho_I > \rho_G$ so that the idiosyncratic component is more persistent than the aggregate component.

Household log income $\ln Y_{i,t}$ is the sum of independent AR(1) processes, so $\ln Y_{i,t}$ is an ARMA(2,1). Appendix A derives the parameters of this composite time series.

The household may hold a risk-free asset A_t which pays exogenous interest rate r. The household's budget constraint is

$$Y_{i,t} + (1+r)A_{i,t} = C_{i,t} + A_{i,t+1} \tag{4}$$

with $A_{i,t+1} \geq 0$ for $t \geq 0$. This implies that households cannot borrow.

2.2 The Information Friction

Households do not observe the incomes or choices of any other households, nor of the aggregate economy. They observe their income $Y_{i,t}$ but cannot independently observe the idiosyncratic and aggregate components $Y^I_{i,t}$ and Y^G_t. Thus if their income rises, they are unsure to what extent the increase was specific to them or economy-wide. Formally, the household's information set evolves by

$$\Omega_{i,t} = \{\Omega_{i,t-1}, Y_{i,t}, A_{i,t}\} \tag{5}$$

The information friction makes households over-estimate the persistence of an aggregate income shock. The autocorrelation of the aggregate income component ρ_G is less than that of idiosyncratic income component ρ_I. The sum of the two components, which is observed by households, has an autocorrelation between ρ_G and ρ_I: individual income is more persistent than aggregate income. So when there is an aggregate income shock and households cannot tell that the shock is aggregate, they expect their income to change more persistently than they would if they had full information.

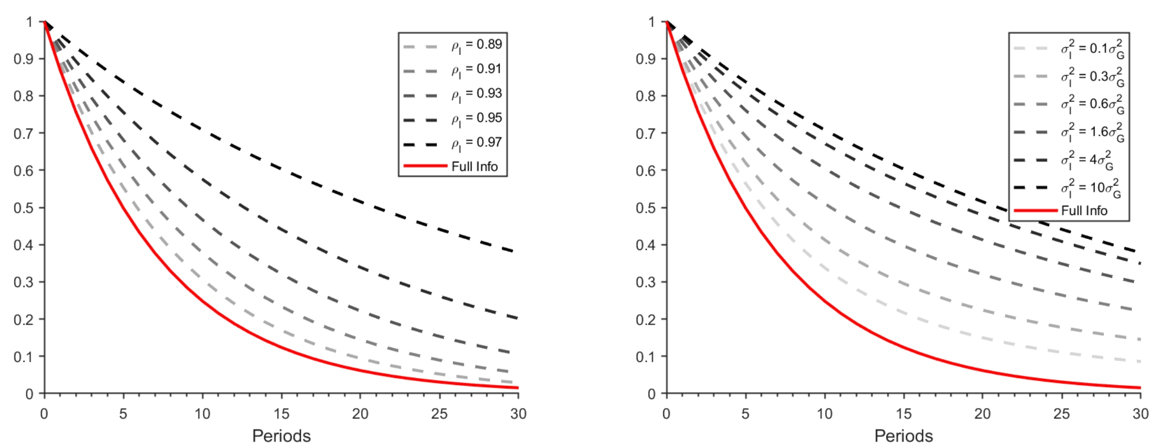

(a) Different Idiosyncratic Persistences

(b) Different Idiosyncratic Variances

Figure 1: Income Forecasts After Aggregate Shocks

Figure 1 demonstrates this over-estimation of an aggregate shock's persistence. This figure plots households' forecasts of their income multiple periods into the future after receiving an aggregate income shock.[5] The aggregate autocorrelation is $\rho_G = 0.8$; when households have full information, they correctly forecast their future income which decays relatively rapidly (the solid red curves).

[5] Appendix A.3 derives expressions for these forecasts.

The information friction matters most when the aggregate and idiosyncratic autocorrelations are most dissimilar. Panel (a) plots forecasts under incomplete information for different idiosyncratic autocorrelations ρ_I. When $\rho_I = \rho_G$, the information friction has no effect and households' forecasts are equivalent to the full information forecasts. When ρ_I is larger, individual income is more persistent, so households' forecasts of future income decay more slowly, and their expectations diverge from the full information case.

The information friction also has larger effects when the idiosyncratic shock has a larger variance. The autocorrelation of individual income is somewhere between those of the aggregate and idiosyncratic components, and when the idiosyncratic component has a larger variance, its larger autocorrelation makes individual income more persistent. Panel (b) plots this effect for different idiosyncratic shock variances σ_I^2, while the other parameters are held at $\rho_G = 0.80$, $\rho_I = 0.95$, and the aggregate shock variance is $\sigma_G^2 = 1$. When σ_I^2 is larger, the idiosyncratic process has higher weight in determining income which becomes more persistent, and forecasted income decays more slowly. When σ_I^2 is small, household income is mostly driven by the aggregate process so households forecast accurately after an aggregate shock; as σ_I^2 goes to zero, the effect of the information friction disappears.

2.3 Equilibrium Definition

Given infinite sequences of exogenous variables $\{Y_{i,t}, Y_t^G, Y_{i,t}^I, u_t^G, u_{i,t}^I\}$ for all $i \in \mathcal{I}$, a competitive equilibrium in this economy consists of infinite sequences of allocations $\{C_{i,t}, A_{i,t}\}$ for all $i \in \mathcal{I}$; and information sets $\Omega_{i,t}$ for all $i \in \mathcal{I}$ such that:

1. Households maximize utility (1), subject to the budget constraint (4) and the no-borrowing constraint.

2. Income is determined by (2) and (3).

3. Information sets evolve according to (5).

3 Quantitative Analysis

In this section we document the model behavior. The information friction raises consumption volatility and nearly eliminates the correlation between household income and their consumption elasticity to aggregate shocks.

3.1 Calibration

We calibrate the model to match features of U.S. states, which we treat as small open economies. The time frequency is annual, and we set the world annual real interest to 2%. The discount factor β is set to be 0.946 in order to match the U.S. ratio of net worth to labor income in the National Income and Product Accounts (NIPA) and the Federal Reserve's Flow of Funds Tables. The risk aversion parameter γ is equal to 1 so that agents have a log utility function. To parameterize the stochastic process for idiosyncratic income, we set the values of ρ_I and σ_I to match the income dynamics estimated by Guvenen et al. (2021), which implies $\rho_I = 0.97$ and $\sigma_I = 0.19$. For the aggregate income shock process, we use the U.S. NIPA accounts to estimate a state-level aggregate income process (Appendix C). We find that the autocorrelation is approximately 0.87 and the standard deviation is roughly 0.03. These statistics define the values for ρ_G and σ_G. Table 1 summarizes our baseline calibration, although we explore some alternative parameter values in Section 4.1.

Table 1: Calibration

Parameter	Interpretation	Value	Reference
β	Discount factor	0.946	U.S. net-worth-to-earnings ratio of 8
r	Real interest rate	0.02	Standard value
γ	Risk aversion	1	Standard value
ρ_I	Persistence of idiosyncratic income shock	0.97	Guvenen et al. (2021)
σ_I	Standard deviation of idiosyncratic income shock	0.19	Guvenen et al. (2021)
ρ_G	Persistence of aggregate income shock	0.87	NIPA
σ_G	Standard deviation of aggregate income shock	0.03	NIPA

The discretization of the stochastic processes for idiosyncratic and aggregate income is different for the full information and the incomplete cases. As shown in Appendix A, the income process in the incomplete information case follows an ARMA(2,1) process. We express this process as a VAR(1) and then use Tauchen (1986)'s approach. We solve the model using a variation of Coleman (1990)'s time iteration method. See Appendix B for more details about the solution method or grids. We consider an asset grid with a zero lower bound for assets, so that in the baseline calibration agents are not able to borrow.

3.2 Main Results

3.2.1 Long-run Moments

In order to assess the role of incomplete information we compare the two models in terms of their long-run moments, and in the way agents respond to idiosyncratic and aggregate income shocks. Using the policy functions we simulate an economy composed of 2,000 individuals for 10,000 periods. We simulate this economy with both incomplete and full information, employing the same sequences of shocks in each simulation. Table 2 presents a summary of the long-run moments of aggregate and cross-sectional moments of the two models.

Table 2: Long-run Moments

	Full Information	Incomplete Information
Aggregate Dynamics		
Consumption: Standard Deviation (log change)	0.0079	0.0104
Consumption: Autocorrelation	0.979	0.965
Assets: Standard Deviation (log change)	0.0070	0.0040
Assets: Autocorrelation	0.998	0.997
Cross-Sectional Statistics		
Income: Mean	1.38	1.38
Income: Coefficient of Variation	0.95	0.95
Consumption: Mean	1.54	1.61
Consumption: Coefficient of Variation	0.81	0.80
Assets: Mean	7.73	11.09
Assets: Coefficient of Variation	1.46	1.12

Notes: Long-run moments are calculated from a simulation of 2,000 households and 10,000 periods. We use the same sequences of aggregate and idiosyncratic shocks for both models.

The simulated statistics of Table 2 reveal how the information friction distorts consumption decisions. Aggregate consumption growth is 32% more volatile under incomplete information, because households undersave in response to aggregate income shocks, which they cannot distinguish from more persistent idiosyncratic shocks. Consumption is also less autocorrelated, reflecting that households are less effective at smoothing consumption. Because the incomplete information households forecast income less accurately, they have a stronger precautionary savings motive. Facing the same interest rate, they hold more assets than they would under full information. The additional financial income allows them to afford higher average consumption as well.

In the cross-section, the information friction distorts consumption and assets in different ways. The friction decreases wealth inequality because the increased precautionary savings motive is strongest at lower asset levels: poor households have stronger incentives to save and move away from the constraint, but rich households still act as if they are nearly unaffected by the constraint. This effect is clear in Figure 2 panel (a), which presents the ergodic distributions of aggregate assets for both models. The information friction distorts the distribution most for low asset levels: the full information model has much more mass near the borrowing constraint, but a similar right tail.

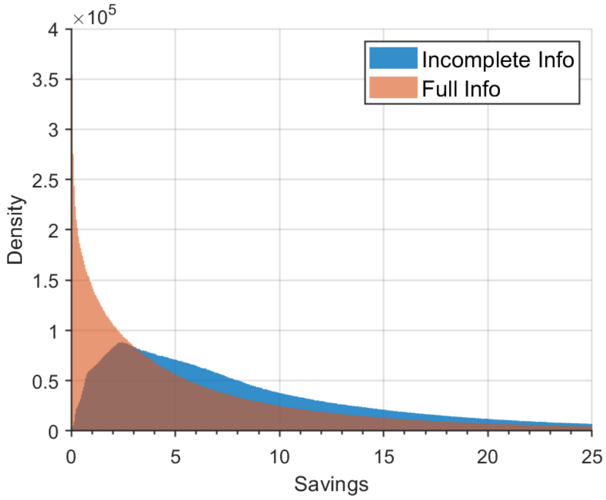

(a) Assets

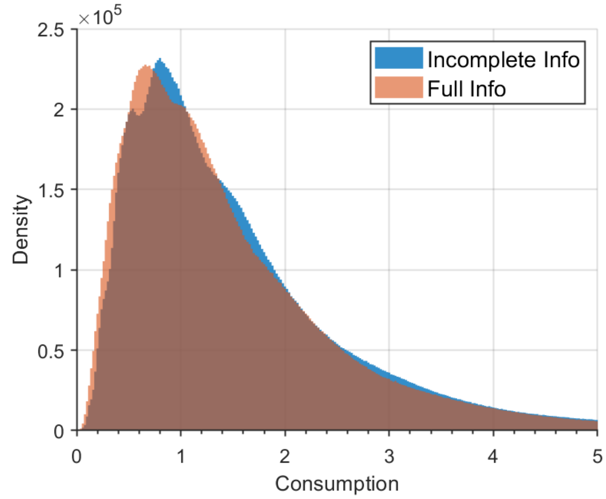

(b) Consumption

Figure 2: Ergodic Distributions - Full and Incomplete Information Models

Notes: The ergodic distributions are each calculated from a simulation of 2,000 households and 10,000 periods. We use the same sequences of aggregate and idiosyncratic shocks for both models. Both distributions extend outside the axis range, but the right tails are omitted for readability.

In contrast, the friction has little effect on consumption inequality (Figure 2 panel (b).) All else equal, the lower wealth inequality would reduces consumption. But this force is offset because households are less effective at consumption smoothing. Most of their income is driven by idiosyncratic shocks, to which households oversave in the short run, before appropriately increasing their consumption response once they realize that their income change was persistent. This delayed consumption response amplifies consumption dispersion because households with large shocks have additional oversavings to draw down as excess consumption. On net, the coefficient of variation for consumption is almost as large with the information friction as it is without, even though wealth is much more equally distributed.

3.2.2 Consumption Volatility

Table 2 reports that the information friction increases aggregate consumption volatility. To understand why, this section compares how the two economies respond to income shocks.

Figure 3 present the impulse response functions of aggregate consumption and assets to an aggregate income shock. The income shock is a one standard deviation increase in u_t^G. In order to calculate the impulse response functions we follow a procedure similar to the one used in Gilchrist et al. (2014). We simulate an economy where the aggregate shock is set to its long-run average for 400 periods. Then we shock the economy at period 401 and compare it to a counterfactual economy where the aggregate shock remains at its long-run average.

The difference in responses is what we report as the impulse response function.

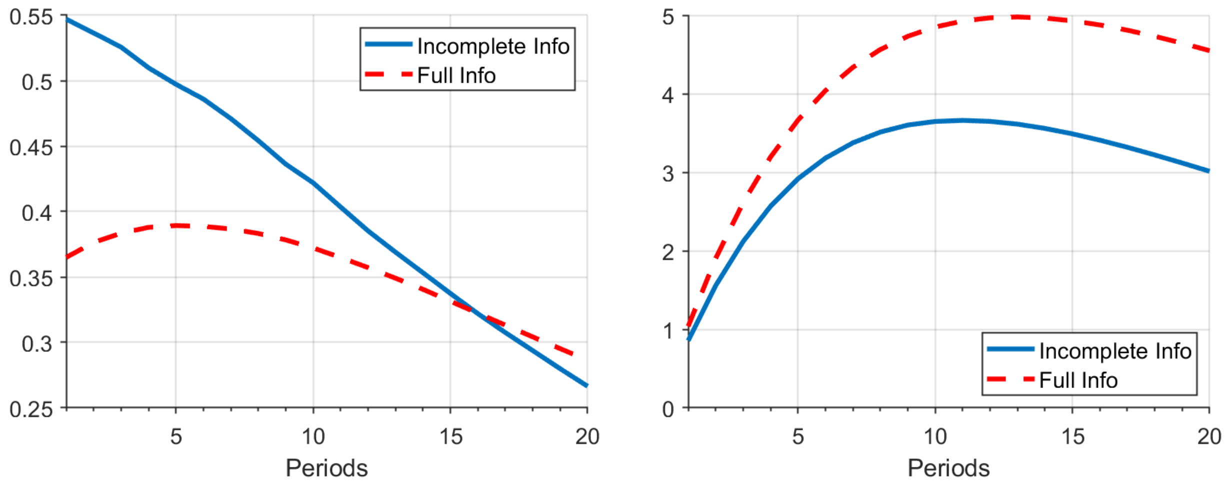

(a) Consumption

(b) Assets

Figure 3: Impulse Responses to an Aggregate Income Shock

Notes: Impulse response functions are calculated by subjecting the economy to a one standard deviation aggregate income shock, and comparing with a counterfactual economy receiving no shock. The impulse response functions are reported as the difference in consumption or assets, normalized by the size of shock.

The responses of consumption and assets differ substantially across models. The response of aggregate consumption, on impact, is nearly 50% larger in magnitude when agents face information frictions. For assets we see the opposite behavior: agents save more in the full information setting. The full information case is the standard response that we would expect to see when agents react to transitory income shocks: upon receiving the income shock, consumption should increase modestly while most of the additional income should be saved. In the presence of financial frictions this behavior is partially mitigated since agents near the borrowing constraint consume a large fraction of the additional income. Under incomplete information, agents cannot initially distinguish if the income shock they are experiencing is an aggregate shock or the more persistent idiosyncratic shock, so their consumption responds much more, and their savings responds less. In short, under incomplete information agents tend to *undersave* in response to aggregate shocks.

We also compute the responses to idiosyncratic income shocks. We follow a similar procedure to the one employed to generate impulse responses to aggregate income shocks. The idiosyncratic income shock consists of a one standard deviation increase to $u^I_{i,t}$. Figure 4 presents the impulse response functions to the idiosyncratic income shock.

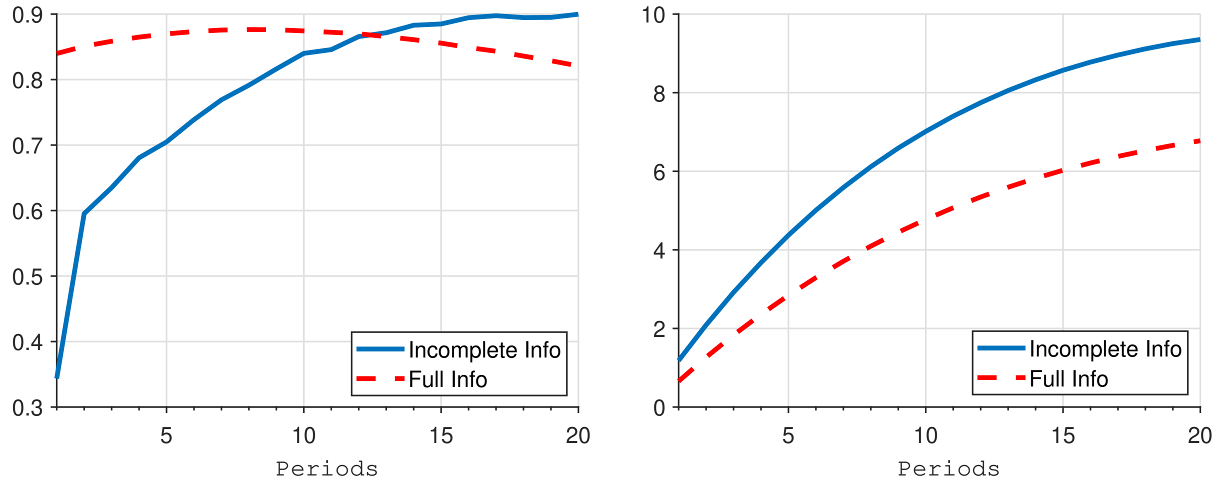

(a) Consumption (b) Assets

Figure 4: Impulse Responses to an Idiosyncratic Income Shock

Notes: Impulse response functions are calculated by subjecting a households to a one standard deviation idiosyncratic income shock, and comparing with a counterfactual household receiving no shock. The impulse response functions are reported as the difference in consumption or assets, normalized by the size of shock.

Figure 4 shows that the relative responses to idiosyncratic income shocks have reversed across models: under incomplete information agents tend to save *more* of the income shock than under the full information scenario. The full information response is consistent with what we would expect to observe under the permanent income hypothesis, whereby "permanent" (in this case very persistent) increases in income translate nearly one-for-one into higher consumption. Since agents cannot tell whether their shock is aggregate or idiosyncratic on impact, they tend to *oversave* in response to idiosyncratic shocks.

The economy features heterogeneous oversaving and undersaving across households. To demonstrate, we calculate $CIE_{i,t}$, the consumption-income elasticity (CIE) of household i in period t:

$$CIE_{i,t} = \frac{\log(C_{i,t}) - \log(C_{i,t-1})}{\log(Y_{i,t}) - \log(Y_{i,t-1})} \tag{6}$$

We study the CIE rather than the well-known MPC in order to directly compare with GOP's evidence. Figure 5 presents the ergodic distribution of CIEs under incomplete and full information. The incomplete information economy features smaller CIEs on average because agents in this economy are worse forecasters and thus have a stronger precautionary savings motive. Idiosyncratic shocks drive most income changes, and agents with full information can immediately observe that these shocks have persistent effects, so they change consumption more elastically than the incomplete information agents. Some CIEs are negative because agents may see their income increase, but by less than they expect, so they reduce consumption in response. This is less common under full information, where the AR(1) structure makes such events less likely. The full information distribution also has larger mass at one, because full information households are much more likely to be borrowing constrained, as shown in Figure 2.

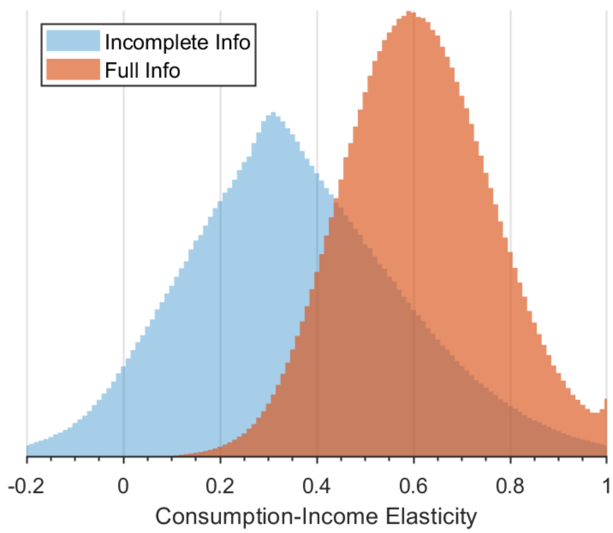

Figure 5: Ergodic Distributions of Consumption-Income Elasticities

Notes: The ergodic distributions of household-level CIEs are each calculated from a simulation of 2,000 households and 10,000 periods. We use the same sequences of aggregate and idiosyncratic shocks for both models. CIEs in the plotted distributions are only included for households experiencing income changes of one standard deviation or more in absolute value.

Why is aggregate consumption so much more volatile when information is incomplete if the CIEs are lower than under full information? Crucially, the CIEs are different in response to idiosyncratic versus aggregate income changes. When we distinguish the generic CIE from the consumption elasticity to *aggregate* income shocks, we find that the incomplete information households are *more* elastic to aggregate income changes. We explore this distinction in the next section.

3.2.3 Elevation and Homogenization of Consumption Elasticities to Aggregate Shocks

Guntin et al. (2020) find that the elasticity of consumption to aggregate shocks is both large and homogeneous across the income distribution. In effect, households respond to transitory aggregate shocks as if they perceive them to be permanent. This is exactly how the information friction in our model affects households. Thus, we find that introducing the information friction to a heterogeneous agent model *elevates and homogenizes consumption elasticities to aggregate income.*

To characterize this effect, we calculate the consumption-income elasticities to aggregate income. Like GOP, we focus specifically on large aggregate shocks, which in our model affects every agent across the income distribution proportionately.[6] The consumption elasticity to aggregate income Y_t^G for individual i at time t is

$$CIE_{i,t}^G = \frac{\log(C_{i,t}) - \log(C_{i,t-1})}{\log(Y_t^G) - \log(Y_{t-1}^G)} \tag{7}$$

We calculate the elasticities for each agent in our simulation for each information structure. To ascertain the cross-sectional relationships with income and wealth, Figure 6 presents the within-decile averages of $CIE_{i,t}^G$, across models. In both cases deciles are calculated from the ergodic distribution of the incomplete information model, so that levels are comparable across information structures.

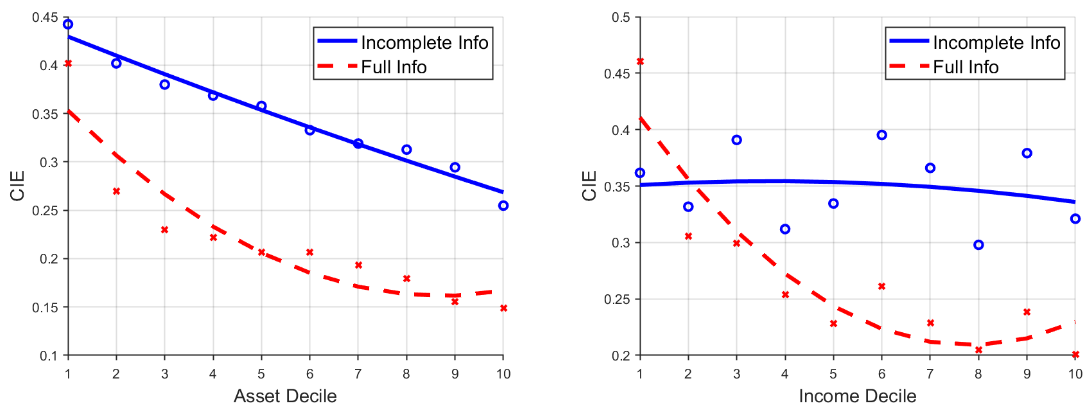

(a) Elasticities by Wealth (b) Elasticities by Income

Figure 6: Consumption-Income Elasticities to Aggregate Income

Notes: The solid and dashed curves are fit from quadratic regressions. The distributions of CIE^G are each calculated from a simulation of 2,000 households and 10,000 periods. We use the same sequences of aggregate and idiosyncratic shocks for both models. The elasticity is calculated at the household level and averaged within household groups corresponding to the asset or income deciles of the incomplete information model's ergodic distribution. Households are grouped based on their position in period $t-1$ for a shock that occurs in period t. The plotted elasticities only include periods with aggregate shocks exceeding two standard deviations in absolute value.

The elasticity of consumption to aggregate income is elevated in the incomplete infor-

[6] We consider aggregate income changes larger than two standard deviations, although our conclusions are not dependent on this particular threshold. In addition to following GOP, setting a threshold for income changes prevents us from occasionally calculating excessively large CIE's when the denominator happens to be small.

mation model, where the average CIE^G is 0.35, versus 0.27 for full information. Figure 6 panel (a) plots the average elasticities within asset deciles, where the information friction (blue circles, with solid blue quadratic fit) substantially elevates the elasticities relative to full information (red crosses, with dashed red quadratic fit) across most of the asset distribution. The elasticities are only similar at very low levels of wealth, where agents have high elasticities because they are likely to be constrained. How can this relationship be so different from the general CIEs (Figure 5) which were much higher under full information? The full information households are extremely elastic to idiosyncratic shocks which drive the majority of income changes, but less elastic to aggregate income changes which are much less persistent. However incomplete information households have similar elasticities to both types of shocks, because they cannot distinguish between them.

The information friction's homogenization effect is clear in the relationship between the consumption elasticity and household income (Figure 6 panel (b)). Homogenization occurs in panel (a) as well, but we focus on the relationship with income in order to mirror the findings by GOP: the consumption elasticity is homogeneously large across the income distribution, in contrast to the negative relationship implied by full information. Why does this occur? Under full information, the response is heterogeneous due to the financial friction; low wealth individuals are more elastic to transitory income shocks because they are near the borrowing constraint. However, the response becomes more homogeneous as income becomes more persistent; in the extreme case when all income shocks are permanent, all agents have the same unit elasticity. Thus under incomplete information where households perceive aggregate shocks as more persistent than they really are, they react more homogeneously.

Our results differ from GOP's findings in two ways. First, we show that the information friction elevates average CIE^G, but not to levels as large as GOP's estimates $(0.7 - 1.2)$. This is because agents in the model mistake aggregate income for idiosyncratic income, which in the conservative baseline calibration only has autocorrelation $\rho_I = 0.97$. If idiosyncratic income were more persistent, then the average CIE^G would be even larger. Second, in some countries GOP find that the relationship between the consumption elasticity and household income is distorted so much as to be upward-sloping in income. This is possible in our model for some alternative calibrations. In particular, when idiosyncratic income has a higher autocorrelation ρ_I, households mistakenly perceive aggregate shocks to be nearly permanent, which elevates elasticities enough to be increasing with income. We study this case in Section 4.1.2.

4 Interactions Between the Frictions

The main purpose of introducing dispersed information into a heterogeneous agents framework was to understand how the information friction affect distributions, particularly the elevation and homogenization documented by GOP, but also the general patterns of consumption and wealth inequality that we discuss in Section 3.2.

However, the information friction and the financial friction interact in rich ways. For example, in this section we show that the information friction attenuates the effects of the financial friction on the wealth distribution, and on the dynamics of inequality. But we also show that the information friction reverses the effects of idiosyncratic risk on the aggregate consumption elasticity, and introduces new cross-sectional heterogeneity, skewness, and correlations for household forecasts. These substantial interactions are further motivation to study incomplete information and heterogeneous agents in a unified framework.

4.1 Parameter Sensitivity

4.1.1 The Borrowing Constraint

To ascertain the impact of the financial friction on our economy, we solve the model for several values of the borrowing constraint, ranging from the no-borrowing baseline to the natural borrowing limit.[7]

Figure 7 plots asset distributions under incomplete and full information for each of these two extreme values. In both cases, reducing the lower bound on assets weakens the distortion that precautionary savings has on the asset distribution: the distribution shifts left of zero, with large masses of agents borrowing to smooth income shocks.

However, this effect is asymmetric across information structures. Under incomplete information, households have a strong precautionary savings motive because they are poor forecasters; few of them choose to be constrained even when unable to borrow, and so relaxing the constraint has little effect on the distribution of assets. But full information households have a weaker precautionary savings motive, often choosing to go to the constraint. When the constraint is relaxed, many more households borrow than under incomplete information.

Therefore, we conclude that the information friction interacts to attenuate the financial friction. By raising the precautionary savings motive, the information friction makes the borrowing constraint less distortionary.

[7] The natural borrowing limit is $-\bar{A} = y_{\min}/r$, where $y_{\min}$ denotes the lowest possible value in the income grid, which in logs is -4.5 standard deviations. We construct the income grids to imply the same natural borrowing limit for both models.

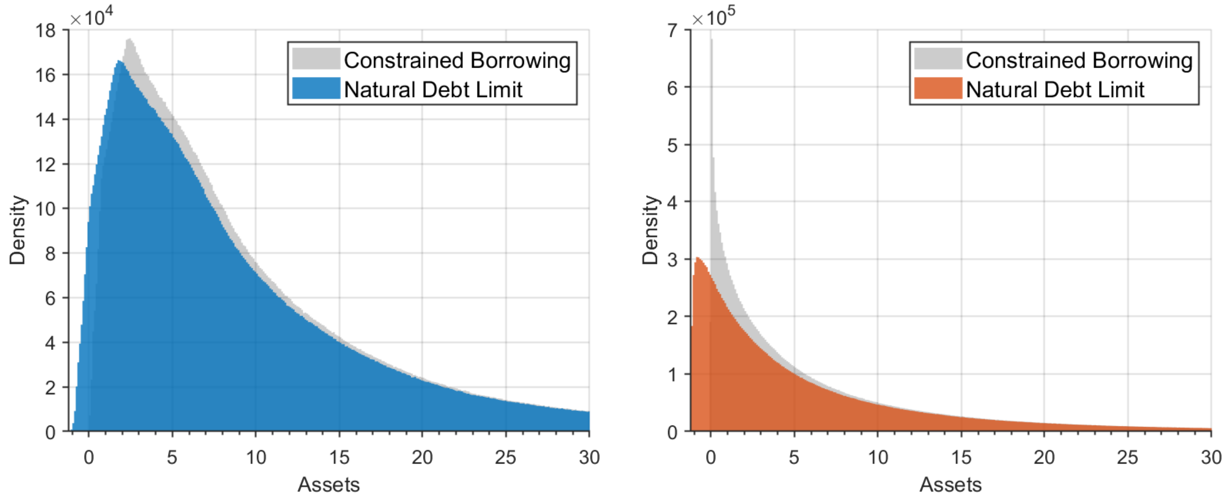

(a) Incomplete Information (b) Full Information

Figure 7: Ergodic Distribution of Assets with the Natural Borrowing Limit

Notes: The ergodic distributions are each calculated from a simulation of 2,000 households and 10,000 periods, experiencing the same sequences of aggregate and idiosyncratic shocks. The light gray distributions are the corresponding ergodic distributions when no borrowing is allowed (Figure 2).

4.1.2 Idiosyncratic Risk

How does the economy change when we adjust the dynamics of idiosyncratic income risk? To address this question, we consider alternative values of ρ_I, the autocorrelation on the idiosyncratic component of income Y^I. This parameter has no effect on how full information agents forecast aggregate income, but under incomplete information, increasing ρ_I makes forecasting aggregate income more difficult by making the combined income process more persistent (as demonstrated in Figure 1.) Figure 8 plots several summary statistics for a range of values of ρ_I.

When idiosyncratic income is more persistent, households increasingly mistake aggregate shocks for permanent income shocks. This strengthens the main mechanisms of our model. Increasing ρ_I raises aggregate consumption volatility (Figure 1 panel (a)) and raises consumption elasticities to aggregate income (panel (b)). Increasing ρ_I also increases the homogenization effect documented by GOP: the slope of the CIE^G-income relationship rises (panel (c)). If idiosyncratic income is sufficiently persistent, the slope can become positive. GOP document a positive relationship for several countries, but this never occurs in our full information model.

The autocorrelation ρ_I has a nonmonotonic effect on the precautionary savings motive. Panel (d) makes this clear, plotting the ratio of average wealth to average income. At most levels, increasing ρ_I increases household income risk. However, when idiosyncratic income

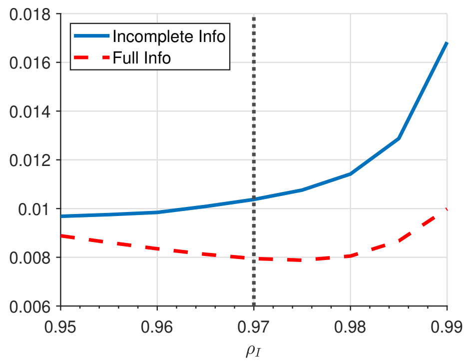

(a) Consumption Volatility

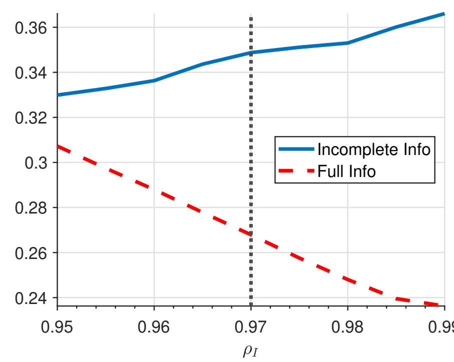

(b) Aggregate CIE^G

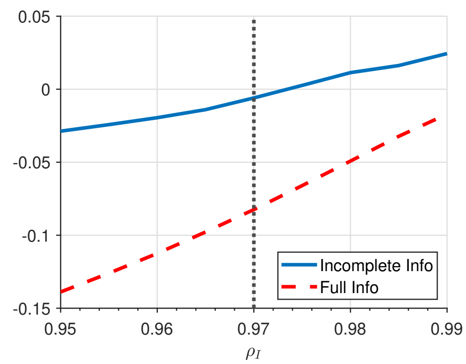

(c) Slope of CIE^G-Income Curve

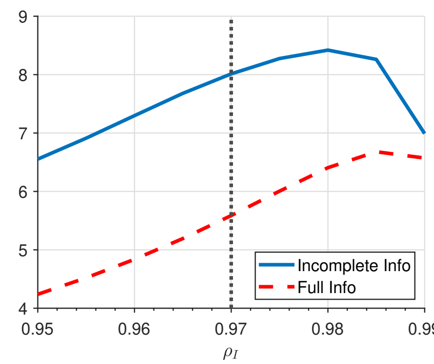

(d) Wealth-Income Ratio

Figure 8: Sensitivity to Idiosyncratic Persistence

Notes: We consider the same sequence of shocks for both models, for every possible value of ρ_I. The baseline $\rho_I = 0.97$ is marked with a dotted line in each panel. Each statistic is calculated from a simulation of 2,000 households and 10,000 periods, experiencing the same sequences of aggregate and idiosyncratic shocks.

becomes extremely persistent, the precautionary savings motive declines. For intuition, consider the limit: when income shocks are completely permanent, there is no precautionary savings motive at all, because consumption follows income one-for-one.

These effects are not common across information structures. The full information CIE^G is decreasing in ρ_I , while it rises under incomplete information. Why? Under full information, the CIE^G moves inversely to the wealth-income ratio, because households consume more when they hold greater wealth, so their consumption is less elastic to aggregate income shocks. The information friction breaks this relationship: when information is incomplete, ρ_I monotonically increases the CIE^G. Under both information structures, a larger ρ_I makes consumption more elastic to idiosyncratic income shocks. But with the friction, households cannot distinguish aggregate from idiosyncratic shocks, so their consumption choice must be more elastic to both types of shocks.

4.2 Inequality Dynamics

One valuable feature of heterogeneous agent models is the ability to study the dynamics of inequality. Introducing the information friction changes these dynamics in nontrivial ways. To demonstrate these effects, Figure 9 plots the average response of inequality measures to a one standard deviation aggregate income shock u_t^G. We calculate inequality as the standard deviation of logs, and for this exercise alone we consider "assets" as cash-on-hand (i.e. financial assets plus current income) so that borrowing constrained households do not have undefined log assets.

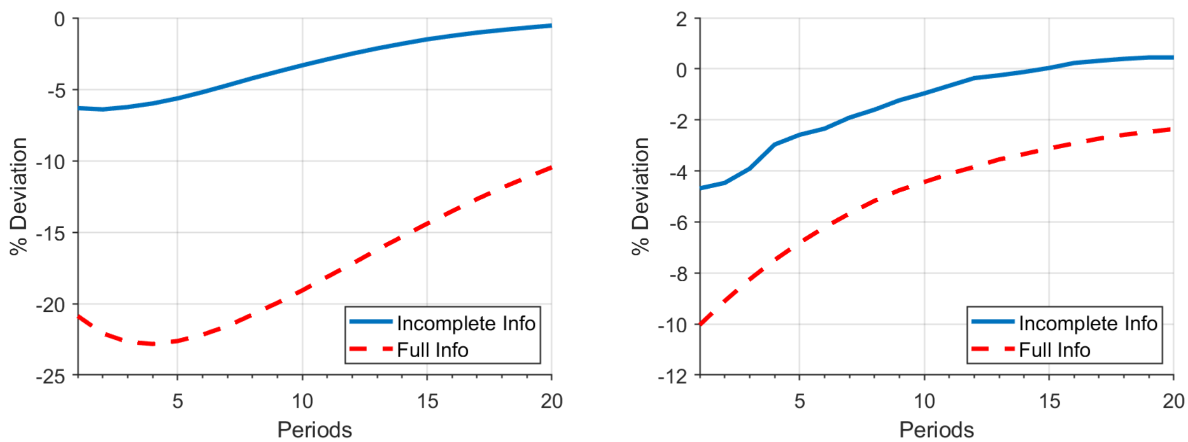

(a) Std. Deviation of Log Assets

(b) Std. Deviation of Log Consumption

Figure 9: Inequality Response to Aggregate Shocks

Notes: Impulse response functions are calculated by subjecting the economy to a one standard deviation aggregate income shock, and comparing with a counterfactual economy receiving no shock. The impulse response functions are reported as the percentage point difference in the standard deviation of log consumption or log assets, relative to the counterfactual economy, and normalized by the size of shock.

Under full information, the standard model predicts that a positive aggregate income shock should reduce consumption and asset inequality (Figure 9, dashed lines). This reduction occurs because all incomes increase proportionately, and income is distributed more equally than assets. To understand this effect, it is useful to view a household's "total wealth" as the sum of financial wealth (i.e. the assets in the model) and human capital (i.e. the present value of future income $Y_{i,t}$) because, absent any financial friction, total wealth would entirely determine consumption. The aggregate shock reduces the share of households' total wealth that is held as financial assets and increases the share held as human capital. As usual, financial assets are distributed more unequally than income, so shifting towards human capital reduces consumption inequality (panel (a)). Similarly, the shift towards human capital causes savings to be distributed more equally, reducing asset inequality (panel (b)).

Under incomplete information, agents have a stronger precautionary savings motive, so they hold more financial assets. Therefore when the aggregate shock increases incomes, it has a smaller effect on the shares of total wealth held as financial assets and human capital. The shock induces a smaller shift towards the more equally distributed human capital than under full information, attenuating the reductions in consumption and asset inequality.

4.3 Forecast Heterogeneity

One implication of heterogeneity among agents is that there is heterogeneity of forecasts. This is true of any model with a persistent income process. But the information friction introduces an additional complication: there is heterogeneity of forecasts *about aggregate variables.*

There is clear empirical evidence that households have heterogeneous forecasts about the macroeconomy.[8] This heterogeneity requires information frictions because FIRE agents all form the same expectations. But there is an additional interaction between the friction and the agent heterogeneity: in linear dispersed information models, the average agent typically holds the average expectation, so the heterogeneity of expectations is irrelevant for macroeconomic dynamics. A consequence of the heterogeneous agent framework is that forecasts about aggregates are nonlinearly related to the endogenous distributions of wealth and consumption.

What is the mechanism? Income is persistent, so higher income individuals expect higher income in the future. Because they cannot disentangle idiosyncratic from aggregate incomes, agents that have higher forecasts of their own income also have higher forecasts of aggregate income. This relationship is strictly mechanical, following from the assumed income process.[9] But income is endogenously correlated with wealth and consumption, so forecasts of aggregates are endogenously correlated as well.

Figure 10 plots the joint distributions of household forecasts of aggregate log incomes and other quantities in the incomplete information model. The joint distribution with income is plotted in panel (a): this relationship is mechanical, entirely implied by the assumed income process and information friction. When households receive higher income, they tend to save, so wealth is positively correlated with income and thus the forecast in panel (b). However, households with higher income do not save it all; they also consume, which thus is positively correlated with forecasts in panel (c).

[8] A large literature documents how heterogeneous forecasts about macroeconomic variables are correlated with household decisions, including Vissing-Jorgensen (2003), Egan et al. (2021a), Egan et al. (2021b), Coibion et al. (2021), and Coibion et al. (2022).

[9] Appendix A.4 derives this relationship.

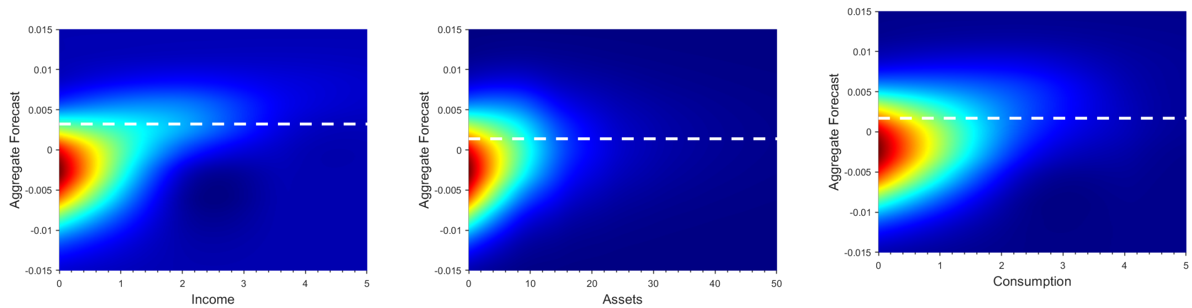

(a) Correlation with Income (b) Correlation with Assets (c) Correlation with Consumption

Figure 10: Distributions of Expectations

Notes: The heatmaps display the ergodic joint distributions in the incomplete information model of: (1.) household forecasts of aggregate income and (2.) either income, assets, or consumption. Red regions indicate the highest density, while blue regions indicate the lowest. The dashed white lines mark the average household, weighted by the x-axis quantity.

The joint distributions have three common patterns. First, the forecasts feature substantial heterogeneity. Second, optimism about the aggregate economy is positively correlated with income, wealth, and consumption. Third, the joint distributions are all skewed with respect to the x-axis. This inequality is typical in heterogeneous agent models. But it has a crucial interaction with the information friction: the skewness biases any weighted-average of forecasts. This can be seen in Figure 10, where the dashed white lines plot the average forecast, weighted by the corresponding x-axis variable. In all cases, this weighted average is greater than the unweighted average, which is necessarily zero.

Our model is simple, but we expect these patterns hold in more general settings so long as income is sufficiently correlated with wealth and consumption. In other models, the consequences of these patterns depend on what matters for the macroeconomy: for example, if it is the forecast associated with the average asset (rather than the average household) that matters for the macroeconomy, then this unequal joint distribution can further distort aggregate dynamics.

5 Corroborating Evidence for the Mechanism

In this section we test whether household expectations of their future income respond differently to idiosyncratic and aggregate shocks.

5.1 Data

To document the response of household income forecasts to idiosyncratic and aggregate shocks, we employ data from the New York Fed Survey of Consumer Expectations (SCE). The SCE is a monthly survey of aggregate and household-level economic conditions and forecasts. It consists of a nationally representative rotating panel of approximately 1300 American households, which remain in the sample for up to 12 months. The survey has been administered since 2013. For our purposes, we require data on both expected and realized household earnings. Unfortunately, this pair is reported only in the auxiliary labor market module of the survey, which is administered to participants every 4 months.

We primarily use two questions from the labor market module. First, to measure household expectations of future earnings, we use the household head's forecast of their 4-month-ahead earnings. This measure is the answer to:

> *What do you believe your annual earnings will be in 4 months?*
> __________________________________*dollars per year*

which we interpret as the household's forecast of instantaneous annualized earnings four months into the future. The advantage of this measure relative to the earnings forecasts in the general SCE survey, is that it is unconditional. The general SCE asks respondents to forecast their earnings over the following year, but do so conditional on holding a job. We prefer to use an unconditional earnings process, which both fits the model and corresponds to the income process that we estimate in the aggregate.

Second, to measure realized income, we use the household head's current annualized reported earnings. This measure is the answer to:

> *How much do you make* ***before*** *taxes and other deductions at your [main/current] job, on an annual basis? Please include any bonuses, overtime pay, tips or commissions.*
> ____*dollars per year*

measured as gross wages or salaries, which respondents are more likely to report accurately. This question is not asked to individuals who are unemployed or out of the labor force, to whom we assign zero earnings. The earnings expectation question was added only in March 2015, and the most recently released data are for March 2021, giving us 19 time periods. The data set contains 10,080 unique households, appearing on average in 2.2 editions of the module. But not all respondents answer all questions; we are left with 12,332 observations with sufficient data.

The general SCE survey contains additional household-level descriptors. Crucially, we observe the state where respondents reside, so that we can connect households to state-level shocks. We use additional descriptors as controls: we observe the industry in which the head of household either works or was employed most recently, we observe their age and education, and we observe demographic characteristics including ethnicity and gender.

5.2 Regressions

We divide a household's log real labor earnings $y_{i,s,t}$ into an idiosyncratic component $y^{Idio}_{i,s,t}$ and an aggregate component $y^{Aggr}_{s,t}$:

$$y_{i,s,t} = y^{Idio}_{i,s,t} + y^{Aggr}_{s,t}$$

where i indexes households, s indexes their state, and t indexes the 4 month time period. The aggregate component $y^{Aggr}_{s,t}$ is the mean earnings in state s and period t. We aggregate at the state level for several reasons. First, using states rather than the entire US economy provides considerably more observations, which is essential given the short history of the SCE. Second, state-level income is more volatile than aggregate income, which gives our analysis additional power. Finally, we treat states as small open economies, which matches the structure of our model and the motivating evidence from Guntin et al. (2020).

Our main regression estimates how household forecasts depend on income:

$$f_{i,s,t} = \beta^{Idio} y^{Idio}_{i,s,t} + \beta^{Aggr} y^{Aggr}_{s,t} + X_{i,s,t} + \varepsilon_{i,s,t} \tag{8}$$

where i indexes households, s indexes their state, and t indexes the 4 month time period. $f_{i,s,t}$ is the household-level forecast of their 4-month-ahead earnings, $y^{Idio}_{i,s,t}$ and $y^{Aggr}_{s,t}$ are the realized aggregate and idiosyncratic earnings components, and $X_{i,s,t}$ is a vector of household-level controls.

If households have full information, then their forecasts would be given by

$$\text{[FIRE]:} \qquad f_{i,s,t} = E_t[y^{Idio}_{i,s,t+1}] + E_t[y^{Aggr}_{s,t+1}]$$

In the model, the earnings components are both AR(1). And if the income components $y^{Idio}_{i,s,t}$ and $y^{Aggr}_{s,t}$ are each AR(1) with autocorrelation ρ^{Idio} and ρ^{Aggr} respectively, then the forecast (8) under full information would satisfy

$$\text{[FIRE, AR(1)]:} \qquad \beta^{Idio} = \rho^{Idio} \qquad \beta^{Aggr} = \rho^{Aggr} \tag{9}$$

However, if households are unable to distinguish between aggregate and idiosyncratic earnings components, then the forecast (8) would satisfy

$$\text{[Incomplete Info., AR(1)]:} \qquad \beta^{Idio} = \beta^{Aggr} \tag{10}$$

We test these information structures when estimating regression (8). First, we test whether we can reject the incomplete information assumption, i.e. whether $\beta^{Idio} = \beta^{Aggr}$. This is a useful direct test, but failing to reject the null hypothesis is not a confirmation of information frictions. Therefore we also test whether the full information assumption fails.

In order to test the full information relationship between forecast coefficients (9), we require values for the autocorrelations ρ^{Idio} and ρ^{Aggr}. For the idiosyncratic autocorrelation, we turn to estimates in the literature. Guvenen et al. (2021) use administrative panel data on US workers from the Social Security Administration to estimate earnings dynamics. We adopt $\rho^{Idio} = 0.97$, which is their most conservative (smallest) estimate for models with a Gaussian AR(1) process. This result is consistent with Heathcote et al. (2010) who use data from the PSID and also estimate 0.97 as the autocorrelation of the persistent component of individual earnings. For the aggregate autocorrelation, we estimate an AR(1) process directly using state-level earnings series from the National Accounts, and find $\rho^{Aggr} = 0.87$. Appendix C details this estimation.

Wary of possible attenuation biases, we do not necessarily want to reject full information if the forecast coefficients β^{Idio} and β^{Aggr} are summarily lower than their corresponding autocorrelations. Instead, we test whether β^{Idio} is relatively larger than β^{Aggr}, independent of their absolute scale. Specifically, define the scalar ψ as

$$\psi \equiv \left(\frac{\rho^{Aggr}}{\rho^{Idio}}\right)^{\frac{1}{3}}$$

The $\frac{1}{3}$ exponent transforms the annual autocorrelations to match the thrice-yearly frequency of the data used in our regression. Appropriately transformed, the full information relationship (9) implies that $\psi\beta^{Idio} = \beta^{Aggr}$. Furthermore, if $\psi\beta^{Idio} > \beta^{Aggr}$, then full information fails, but in the opposite direction as implied by our model. Therefore our second statistical test is whether $\psi\beta^{Idio} \geq \beta^{Aggr}$.

5.3 Results

Table 3 presents the results of forecast regression (8). Column (1) is the basic regression with no additional controls. In column (2), we control for state-level effects on expectations. We cannot include state-time controls, because aggregate earnings vary at the state level. Industry fixed effects (3) control for the industry and employment status of the household head. Human capital fixed effects (4) control for education and age. Demographic fixed effects (5) control for gender and ethnicity. Columns (6)-(8) include all fixed effects. Column (7) utilizes the panel dimension to control for lagged earnings; this is our preferred specification because our theory assumes that households cannot distinguish between the idiosyncratic and aggregate components of lagged earnings. Still, we also run the specification in column (8) which includes lagged idiosyncratic and aggregate earnings separately. Finally, column (9) uses household fixed effects.

Table 3: Effects of Log Earnings on Household Earnings Forecasts

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Idio. Log Earnings	0.677	0.676	0.665	0.653	0.666	0.633	0.563	0.563	0.311
	(0.0195)	(0.0195)	(0.0206)	(0.0207)	(0.0200)	(0.0223)	(0.0412)	(0.0412)	(0.0106)
Aggr. Log Earnings	0.800	0.750	0.783	0.764	0.785	0.702	0.634	0.605	0.292
	(0.0388)	(0.0518)	(0.0395)	(0.0402)	(0.0394)	(0.0546)	(0.0533)	(0.0493)	(0.0330)
Lag Log Earnings							0.205		
							(0.0302)		
Lag Idio. Log Earnings								0.203	
								(0.0303)	
Lag Aggr. Log Earnings								0.264	
								(0.0434)	
$\beta^{Idio} = \beta^{Aggr}$ p-value	0.002	0.144	0.003	0.005	0.003	0.187	0.074	0.209	0.543
$\psi\beta^{Idio} \geq \beta^{Aggr}$ p-value	0.000	0.053	0.000	0.001	0.000	0.080	0.021	0.062	0.794
Observations	12332	12332	12332	12332	12323	12323	6406	6406	12332
R^2	0.554	0.556	0.558	0.563	0.558	0.572	0.649	0.650	0.551
State F.E.		X				X	X	X	
Industry F.E.			X			X	X	X	
Human Capital F.E.				X		X	X	X	
Demographic F.E.					X	X	X	X	
Individual F.E.									X

Notes: Standard errors in parentheses, clustered at the state-month level. In all cases, the dependent variable is the household-level log forecast of its 4-month-ahead annualized earnings.

In all cases except column (9), the coefficient on aggregate log earnings exceeds that of idiosyncratic log earnings. Household forecasts of future income are more sensitive to changes in aggregate earnings, even though their idiosyncratic earnings are much more persistent! The test results formalize this conclusion. When we test $\psi\beta^{Idio} \geq \beta^{Aggr}$, we reject the inequality at the 5% level in our preferred specification (7). This means that the relative response of household forecasts to aggregate earnings exceeds what it would be under full information, or if full information failed but in the opposite direction than implied by our model. And even when our statistical test fails to reject, our estimates remain a better fit to the incomplete information model than the full information model. Indeed, when we test the assumption of our model ($\beta^{Idio} = \beta^{Aggr}$), we either fail to reject it, or if we do, it is because households are even more responsive to aggregate income than our friction implies.

Specification (9) demonstrates why we choose not to use panel-level fixed effects: they attenuate the estimated coefficients. This reflects a classic problem; households are in the panel for at most 3 periods, so the panel component is very short. In a traditional dynamic panel regression, short panels bias the coefficients (Nickell, 1981). Our panel regression does not have an explicit dynamic component, but it is not safe from Nickell bias concerns, because the household expectations are forecasts of the panel component. Still, even in this specification, the estimated coefficients are close enough that we fail to reject our information structure.

6 Conclusion

In this paper we introduced dispersed information to an otherwise standard open economy heterogeneous agents model. We demonstrated that the information friction increased consumption volatility and reduced heterogeneity in household's response to aggregate income shocks. Then we documented evidence for our central mechanism in US forecast data.

Our central findings are robust, but there is further work to be done. How would the friction interact with capital accumulation in the model? Or a richer asset market or household risks and decisions? What about closed economies? Regardless, it is clear that households respond to aggregate shocks as if they are more persistent than they actually are, so this type of information friction will be valuable in any setting where consumption volatility matters.

References

Adams-Prassl, Abigail, Teodora Boneva, Marta Golin, and Christopher Rauh, "Perceived Returns to Job Search," IZA Discussion Paper 15307, Institute of Labor Economics (IZA) May 2022.

Aguiar, Mark and Gita Gopinath, "Emerging Market Business Cycles: The Cycle Is the Trend," *Journal of Political Economy*, February 2007, *115* (1), 69–102.

Andrade, Philippe, Olivier Coibion, Erwan Gautier, and Yuriy Gorodnichenko, "No firm is an island? How industry conditions shape firms' expectations," *Journal of Monetary Economics*, January 2022, *125*, 40–56.

Angeletos, G. M. and C. Lian, "Chapter 14 - Incomplete Information in Macroeconomics: Accommodating Frictions in Coordination," in John B. Taylor and Harald Uhlig, ed., *Handbook of Macroeconomics*, Vol. 2, Elsevier, 2016, pp. 1065–1240.

Angeletos, George-Marios and Zhen Huo, "Myopia and Anchoring," *American Economic Review*, April 2021, *111* (4), 1166–1200.

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub, "The Intertemporal Keynesian Cross," September 2018.

_ , _ , and _ , "Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated HANK Model," January 2020.

Barro, Robert J., Xavier Sala-I-Martin, Olivier Jean Blanchard, and Robert E. Hall, "Convergence Across States and Regions," *Brookings Papers on Economic Activity*, 1991, *1991* (1), 107–182.

Bilbiie, Florin O., "The New Keynesian cross," *Journal of Monetary Economics*, October 2020, *114*, 90–108.

Broer, Tobias, Alexandre Kohlhas, Kurt Mitman, and Kathrin Schlafmann, "Information and Wealth Heterogeneity in the Macroeconomy," 2021. Publisher: CEPR Discussion Paper No. DP15934.

_ , Alexandre N. Kohlhas, Kurt Mitman, and Kathrin Schlafmann, "On the possibility of Krusell-Smith Equilibria," *Journal of Economic Dynamics and Control*, August 2022, *141*, 104391.

Carroll, Christopher D, "A Theory of the Consumption Function, With and Without Liquidity Constraints," *Journal of Economic Perspectives*, August 2001, *15* (3), 23–45.

Carroll, Christopher D., "Macroeconomic Expectations of Households and Professional Forecasters*," *The Quarterly Journal of Economics*, February 2003, *118* (1), 269–298.

_ , Edmund Crawley, Jiri Slacalek, Kiichi Tokuoka, and Matthew N. White, "Sticky Expectations and Consumption Dynamics," *American Economic Journal: Macroeconomics*, July 2020, *12* (3), 40–76.

Chetty, Raj, John Friedman, Nathaniel Hendren, Michael Stepner, and The Opportunity Insights Team, "The Economic Impacts of COVID-19: Evidence from a New Public Database Built Using Private Sector Data," Technical Report w27431, National Bureau of Economic Research, Cambridge, MA June 2020.

Coibion, Olivier, Dimitris Georgarakos, Yuriy Gorodnichenko, Geoff Kenny, and Michael Weber, "The Effect of Macroeconomic Uncertainty on Household Spending," March 2021.

_ , Yuriy Gorodnichenko, and Michael Weber, "Monetary Policy Communications and Their Effects on Household Inflation Expectations," *Journal of Political Economy*, June 2022, *130* (6), 1537–1584. Publisher: The University of Chicago Press.

Coleman, John Wilbur, "Solving the Stochastic Growth Model by Policy-Function Iteration," *Journal of Business & Economic Statistics*, January 1990, *8* (1), 27–29.

Egan, Mark, Alexander MacKay, and Hanbin Yang, "Recovering Investor Expectations from Demand for Index Funds," *The Review of Economic Studies*, December 2021, p. rdab086.

Egan, Mark L., Alexander MacKay, and Hanbin Yang, "What Drives Variation in Investor Portfolios? Evidence from Retirement Plans," December 2021.

Gallegos, José-Elías, "HANK beyond FIRE." PhD dissertation, Department of Economics, Stockholm University 2022.

Gilchrist, Simon, Jae W. Sim, and Egon Zakrajšek, "Uncertainty, Financial Frictions, and Investment Dynamics," Working Paper 20038, National Bureau of Economic Research April 2014. Series: Working Paper Series.

Guntin, Rafael, Pablo Ottonello, and Diego Perez, "The Micro Anatomy of Macro Consumption Adjustments," *NBER Working Paper*, 2020.

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song, "What Do Data on Millions of U.S. Workers Reveal About Lifecycle Earnings Dynamics?," *Econometrica*, 2021, *89* (5), 2303–2339.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante, "Quantitative Macroeconomics with Heterogeneous Households," *Annual Review of Economics*, 2009, *1* (1), 319–354. _eprint: https://doi.org/10.1146/annurev.economics.050708.142922.

_ , _ , and _ , "The Macroeconomic Implications of Rising Wage Inequality in the United States," *Journal of Political Economy*, August 2010, *118* (4), 681–722.

Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles, "Household Expenditure and the Income Tax Rebates of 2001," *American Economic Review*, December 2006, *96* (5), 1589–1610.

Kaplan, Greg and Giovanni L. Violante, "A Model of the Consumption Response to Fiscal Stimulus Payments," *Econometrica*, 2014, *82* (4), 1199–1239. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA10528.

_ **and** _ , "Microeconomic Heterogeneity and Macroeconomic Shocks," *Journal of Economic Perspectives*, August 2018, *32* (3), 167–194.

_ **and** _ , "The Marginal Propensity to Consume in Heterogeneous Agent Models," May 2022.

_ **, Benjamin Moll, and Giovanni L. Violante**, "Monetary Policy According to HANK," *American Economic Review*, March 2018, *108* (3), 697–743.

Krueger, D., K. Mitman, and F. Perri, "Chapter 11 - Macroeconomics and Household Heterogeneity," in John B. Taylor and Harald Uhlig, eds., *Handbook of Macroeconomics*, Vol. 2, Elsevier, January 2016, pp. 843–921.

Krusell, Per, Lee E. Ohanian, José-Víctor Ríos-Rull, and Giovanni L. Violante, "Capital-skill Complementarity and Inequality: A Macroeconomic Analysis," *Econometrica*, 2000, *68* (5), 1029–1053.

Kydland, Finn E. and Edward C. Prescott, "Time to Build and Aggregate Fluctuations," *Econometrica*, 1982, *50* (6), 1345–1370.

Lucas, Robert E, "Expectations and the Neutrality of Money," *Journal of economic theory*, 1972, *4* (2), 103–124.

Mankiw, N. Gregory and Ricardo Reis, "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *The Quarterly Journal of Economics*, November 2002, *117* (4), 1295–1328.

Nickell, Stephen, "Biases in Dynamic Models with Fixed Effects," *Econometrica*, 1981, *49* (6), 1417–1426.

Parker, Jonathan A., Nicholas S. Souleles, David S. Johnson, and Robert McClelland, "Consumer Spending and the Economic Stimulus Payments of 2008," *American Economic Review*, October 2013, *103* (6), 2530–2553.

Quadrini, Vincenzo and José-Víctor Ríos-Rull, "Inequality in Macroeconomics," 2015.

Tauchen, George, "Finite state markov-chain approximations to univariate and vector autoregressions," *Economics Letters*, January 1986, *20* (2), 177–181.

Vissing-Jorgensen, Annette, "Perspectives on Behavioral Finance: Does "Irrationality" Disappear with Wealth? Evidence from Expectations and Actions," *NBER Macroeconomics Annual*, January 2003, *18*, 139–194. Publisher: The University of Chicago Press.

A The Sum of Independent Income Processes

In this appendix, we characterize the time series properties of log income in the model. Index households by i and time by t. Household log income $y_{i,t}$ is given by

$$y_{i,t} = y^I_{i,t} + y^G_t$$

where $y^I_{i,t}$ is idiosyncratic and mean zero in the population for all t, while y^G_t is aggregate and common to all households.

A.1 The ARMA(2,1) Representation

The idiosyncratic and aggregate components are AR(1):

$$y^I_{i,t} = \rho_I y^I_{i,t-1} + u^I_{i,t}$$

$$y^G_{i,t} = \rho_G y^G_{i,t-1} + u^G_{i,t}$$

with $u^I_{i,t} \sim N(0, \sigma^2_I)$, $u^G_t \sim N(0, \sigma^2_G)$, $\rho_I \in (0, 1)$, and $\rho_G \in (0, 1)$.

It is helpful to use lag operator notation to define these time series, which become:

$$y^I = L\rho_I y^I + u^I$$

$$y^G = L\rho_G y^G + u^G$$

$$y = L\varrho(L)y + w$$

where ϱ is a lag operator polynomial to be found, and w is a white noise process to be found.

It is well known that the sum of AR(1) processes is ARMA(2,1). The autoregressive coefficients parameters are $\varrho_0 = \rho_I + \rho_G$ and $\varrho_1 = -\rho_I\rho_G$:

$$y - (\rho_I + \rho_G)Ly + \rho_I\rho_G L^2 y = y_I + y_G - (\rho_I + \rho_G)L(y_I + y_G) + \rho_I\rho_G L^2(y_I + y_G)$$

$$= \rho_I L y_I + u_I + \rho_G L y_G + u_G - (\rho_I + \rho_G)L(y_I + y_G) + \rho_I\rho_G L^2(y_I + y_G)$$

$$= u_I + u_G - \rho_G L y_I - \rho_I L y_G + \rho_I\rho_G L^2(y_I + y_G)$$

$$= u_I + u_G - \rho_G L(y_I - \rho_I L y_I) - \rho_I L(y_G - \rho_G L y_G)$$

$$= u_I + u_G - \rho_G L u_I - \rho_I L u_G \equiv z$$

The object z is $MA(1)$. What is the structure of the MA term?

$$cov(z, L^j z) = \begin{cases} var(u_I)(1+\rho_G^2) + var(u_G)(1+\rho_I^2) & j = 0 \\ -var(u_I)\rho_G - var(u_G)\rho_I & j = 1 \\ 0 & j > 1 \end{cases}$$

thus we can write

$$z = w + \theta L w$$

To finish characterizing the ARMA process, we need to know the variance of w, and the value of θ. These quantities are related by two equations:

$$var(z) = var(w) + \theta^2 var(w)$$

$$cov(z, Lz) = \theta var(w)$$

which imply

$$0 = var(w)^2 - var(z)var(w) + cov(z, Lz)^2$$

Apply the quadratic formula and take the larger root:

$$var(w) = \frac{var(z)}{2} + \sqrt{\left(\frac{var(z)}{2}\right)^2 - cov(z, Lz)^2}$$

Then calculate θ by

$$\theta = cov(z, Lz)/var(w)$$

Finally, w can be expressed analytically in terms of the underlying shocks:

$$w = (1+\theta L)^{-1} z = z - \theta L z + \theta^2 L^2 z - ...$$

$$w = (1+\theta L)^{-1}(I - \rho_G L)u_I + (1+\theta L)^{-1}(I - \rho_I L)u_G \tag{11}$$

A.2 The VAR(1) Representation

Stack the variables as such:

$$\boldsymbol{y}_t \equiv \begin{pmatrix} y_t \\ y_{t-1} \\ w_t \end{pmatrix}$$

Then $\boldsymbol{y}_t$ is a VAR(1) with coefficient matrix $B = \begin{pmatrix} \varrho_0 & \varrho_1 & \theta \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$ and innovation Cw_t for $C = \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix}$:

$$\boldsymbol{y}_t = B\boldsymbol{y}_{t-1} + Cw_t$$

A.3 Income Forecasts at Various Horizons

This section derives the term structure of expectations under the information friction, which are is used in Section 2.2.

Income follows the ARMA(2,1) process

$$y_{i,t} = \varrho_0 y_{i,t-1} + \varrho_1 y_{i,t-2} + w_{i,t} + \theta w_{i,t-1}$$

so the one-period-ahead forecast is given by

$$E[y_{i,t+1}|\Omega_{i,t}] = \varrho_0 y_{i,t} + \varrho_1 y_{i,t-1} + \theta w_{i,t}$$

because $E[w_{i,t+1}|\Omega_{i,t}] = 0$.

The two-period-ahead forecast is given by

$$E[y_{i,t+2}|\Omega_{i,t}] = \varrho_0 E[y_{i,t+1}|\Omega_{i,t}] + \varrho_1 y_{i,t}$$

beyond this horizon, the h-period-ahead forecast can be found recursively by

$$E[y_{i,t+h}|\Omega_{i,t}] = \varrho_0 E[y_{i,t+h-1}|\Omega_{i,t}] + \varrho_1 E[y_{i,t+h-2}|\Omega_{i,t}]$$

for $h \geq 2$.

A.4 Forecasts of Aggregate Income

In the model, agents have no need to forecast the aggregate economy; they only need to forecast their own income. However, it is possible to construct the forecasts that agents would make, given the information friction that they face.

Agent i's forecast of aggregate income conditional on their information set $\Omega_{i,t}$ is

$$E[y^G_{t+1}|\Omega_{i,t}] = \rho_G E[y^G_t|\Omega_{i,t}]$$

$$= \rho^G \sum_{j=0}^{\infty} \rho_G^j E[u_{t-j}^G | \Omega_{i,t}]$$

How do agents form expectations over past shocks? Linear backcasting is easily expressed in terms of current and past orthogonal forecast errors $w_{i,t}$. Let W^G denote the lag operator polynomial that gives the aggregate component of $w_{i,t}$ from current and past aggregate shocks. Per equation (11), this polynomial is given by $W^G = \frac{1-\rho_I L}{1+\theta L}$. Let W_j^G denote the j coefficient of this polynomial. Then the backcast is given by:

$$E[u_{t-j}^G | \Omega_{i,t}] = \sum_{k=0}^{j} \frac{cov(u_{t-j}^G, w_{i,t-k})}{var(w_{i,t-k})} w_{i,t-k}$$

$$= \sum_{k=0}^{j} \frac{cov(u_{t-j}^G, W_{j-k}^G u_{t-j}^G)}{var(w_{i,t-k})} w_{i,t-k} = \frac{\sigma_G^2}{var(w_{i,t})} \sum_{k=0}^{j} W_{j-k}^G w_{i,t-k}$$

Plugging in this backcasting formula gives the expression for the aggregate income forecast:

$$E[y_{t+1}^G | \Omega_{i,t}] = \frac{\rho^G \sigma_G^2}{var(w_{i,t})} \sum_{j=0}^{\infty} \sum_{k=0}^{j} \rho_G^j W_{j-k}^G w_{i,t-k}$$

B Solution Method

This section provides a description of the algorithm we use to numerically solve for the equilibrium of the full and incomplete information cases.

We start by discretizing the income processes using the approach of Tauchen (1986). For the full information case we discretize the aggregate and idiosyncratic processes separately (we consider 11 points for each income process). For the incomplete information case we write the income process as a VAR(1) (see A.2 for more details) and then discretize it. Note that the VAR(1) case contains three variables. We allow for 11 points for each variable, so the number of exogenous states is $11 \times 11 \times 11 = 1131$.

The asset grid is discrete and consists of 200 points.[10] We skew the allocation of points in the asset grid in order to have a better coverage over lower asset levels.

After discretizing the income processes, we proceed to our time iteration method, which is similar to the one described in Coleman (1990). We start with a conjecture for the asset holdings policy function, A', defined over the state space (Y, A), where Y represents the vector of exogenous states of each model.[11] The steps of the solution algorithm are the

[10]Our results do not change substantially by increasing the number of asset grid points, as our solution method relies on first order conditions.

[11]Note that this guess corresponds to a matrix with dimensions $N_Y \times N_A$, where N_Y and N_A correspond

following:

1. Start iteration j with a guess for $A'_j(Y,A) \geq -\bar{A}$, where $-\bar{A}$ denotes the borrowing limit. Using this guess construct:

$$C_j(Y,A) = Y + A(1+r) - A'_j(Y,A) \tag{12}$$

   Using (12), compute the discounted expected marginal utility

$$\beta(1+r^*)\mathbb{E}_{Y'|Y}\left[u_j(Y', A'_j(Y,A))\right], \tag{13}$$

   where $u_j(Y,A) = C_j(Y,A)^{-\gamma}$.

2. Assume the borrowing constraint binds. Note that when the constraint binds we have that consumption is $C_{j+1}(Y,A) = Y + A(1+r) + \bar{A}$. We check whether this assumption holds by calculating the residual of the Euler equation:

$$\mathcal{R}(Y,A) = u_{j+1}(Y,A) - \beta R^* \mathbb{E}_{Y'|Y}\left[u_j(Y', A'_j(Y,A))\right]. \tag{14}$$

   If $\mathcal{R}(Y,A) > 0$, we keep the values for $C_{j+1}(Y,A)$. Otherwise, the constraint does not bind for that point of the state space and we discard $C_{j+1}(Y,A)$. We then solve for the value of $C_{j+1}(Y,A)$ that satisfies

$$C_{j+1}(Y,A)^{-\gamma} = \beta R^* \mathbb{E}_{Y'|Y}\left[u_j(Y', A'_j(Y,A))\right]. \tag{15}$$

3. Use the resource constraint to obtain the updated conjecture for asset holdings $A'_{j+1}(Y,A) = Y + A(1+r) - C_{j+1}(Y,A)$.

4. Check for convergence. If $||A'_{j+1}(Y,A) - A'_j(Y,A)|| < \epsilon$, then the problem is solved. Otherwise, discard A'_j and use A'_{j+1} as the new guess for the problem (go back to step 1).

C Aggregate Income Persistence

In this section we estimate the autocorrelation of aggregate log labor earnings.

We measure aggregate log real labor earnings $y^{Aggr}_{s,t}$ at the state level from the National Accounts, for state s and year t. We define labor earnings as the sum of wages and salaries,

to the number of elements in the grid of income states and assets, respectively.

supplements to wages and salaries, and proprietors' income. We deflate by the PCE deflator, to match the procedure used by Guvenen et al. (2021) for idiosyncratic earnings. We estimate the following panel regression:

$$y_{s,t}^{Aggr} = \rho^{Aggr} y_{s,t-1}^{Aggr} + d_s + X_{s,t} + \varepsilon_{i,s,t}$$

where d_s denote fixed state-level controls, while $X_{s,t}$ are further controls that vary with time. We estimate this regression both by OLS and by 2SLS, instrumenting for $y_{s,t-1}^{Aggr}$ with $y_{s,t-2}^{Aggr}$. The IV regressions allows for consistent estimation of ρ^{Aggr} even when the error term $\varepsilon_{i,s,t}$ is MA(1). We take this approach in order to be comparable to Guvenen et al. (2021), who allow individual earnings to contain an i.i.d. transitory term, which is equivalent to letting earnings be ARMA(1,1).

Table 4 reports our estimates of ρ^{Aggr}. Columns (1) report results for the regression with state-specific as the only controls, which allows for the longest sample. We let our trends be state-specific, given the well-known heterogeneity of growth rates across states (Barro et al., 1991). This is our simplest specification and yields the largest estimate, but even in this case it is significantly lower the the idiosyncratic persistence. In specifications (2) and (3) we control for state-level demographics, as changes in worker composition affect average earnings in a predictable way. We include data from the Current Population Survey on age, gender, race, and education. Controlling for these factors, we estimate an autocorrelation of 0.89 or 0.87 depending on whether our controls are additive or interacted with a time trend, respectively. The specifications (4)-(5) use the 2SLS approach in order to allow for transitory i.i.d. shocks to earnings. Allowing for these transitory shocks reduces the estimated magnitudes even further: 0.84 and 0.81 respectively. Lastly, we run our estimation for the 1994-2013 subsample, in order to most directly compare our aggregate estimate with the idiosyncratic autocorrelation estimated from this sample period by Guvenen et al. (2021), which is the value for ρ^{Idio} that we use in our calibration. Columns (6) and (7) present these results, with autocorrelations of 0.80 or 0.56; during this time period, income was generally less persistent than in the broader sample.

Table 4: Aggregate Earnings Autocorrelations

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Lag Real Earnings	0.924	0.887	0.871	0.837	0.814	0.798	0.560
	(0.006)	(0.008)	(0.009)	(0.009)	(0.009)	(0.017)	(0.029)
Observations	4599	2686	2686	2686	2686	1020	1020
R^2	0.991	0.995	0.995	0.995	0.995	0.992	0.992
State Trends	X	X	X	X	X	X	X
Demo. F.E.		X		X		X	
Demo. Trends			X		X		X
Transitory Shocks				X	X	X	X
Sample Period	1929-2020	1962-2020	1962-2020	1962-2020	1962-2020	1994-2013	1994-2013

Notes: Heteroskedasticity-consistent standard errors in parentheses. In all cases the dependent variable is state-level average annual labor earnings.

Our preferred specification is column (5), given that we expect the effects of demographic factors on income to change over time, due to female entry into the labor force, changing attitudes towards race, and rising capital-skill complementarity (Krusell et al., 2000). However, we choose specification (3)'s 0.87 as our baseline calibration for ρ^{Aggr}; smaller values will strengthen the effects of the information friction in our model, and we aim to be conservative in our approach. It might be reasonable to choose a lower value, given the lower estimates for the 1994-2013 time period that informs our idiosyncratic process. But, we are concerned that this shorter time period might be susceptible to Nickell bias attenuating the estimates, so we are wary of selecting the low estimates below 0.80.