Income Expectations Under Information Frictions and Aggregate Consumption*

Sanhitha Jugulum Krishna Kamepalli

June 30, 2026

Abstract

In this paper, we document that households respond almost equally to aggregate and idiosyncratic income shocks, even though the aggregate component of income is much more persistent. This pattern is difficult to reconcile with the standard full-information permanent-income logic, which predicts that more persistent shocks should induce larger consumption responses. A quantitative full-information Huggett model with aggregate risk, disciplined by the estimated income process, predicts 4.4 times greater pass-through of aggregate shocks than is observed in the data. We show that incomplete information about the source of income fluctuations provides a natural resolution: households observe total income as a signal to extract the aggregate component, and therefore under-react to aggregate shocks. Embedding this inference problem in a heterogeneous-agent model aligns the model-implied estimates more closely with the empirical pass-through estimates and implies substantially weaker aggregate consumption responses to aggregate shocks.

*Jugulum: Columbia University, email: sanhitha.jugulum@columbia.edu. Kamepalli: Columbia University, email: krishna.kamepalli@columbia.edu. We thank Hassan Afrouzi, Jennifer La'O, Christian Moser, and audiences in the Economic Fluctuations group and the Econometrics group at Columbia University for many helpful comments.

1 Introduction

A central implication of the permanent-income hypothesis is that consumption should respond more strongly to more persistent income shocks. In standard consumption-savings models, a household that understands its income process maps an observed income change into a change in lifetime resources, and this mapping crucially depends on persistence. This logic underlies a large body of quantitative literature on household heterogeneity and aggregate demand, in which borrowing constraints and the wealth distribution shape the magnitude of consumption responses ((Hall, 1978; McKay and Reis, 2016; Kaplan et al., 2018; Auclert, 2019)). However, the present-value content of the shock remains a key force linking persistence to spending.

This paper documents a sharp empirical puzzle for this benchmark. Using PSID panel data on household income and consumption from 1999-2019, we decompose residual labor income into an aggregate persistent component and an idiosyncratic persistent component, and then estimate how consumption responds to shocks to each. The aggregate component of income is highly persistent, with estimated persistence $\hat{\rho}_A = 0.99$, while the idiosyncratic component is substantially less persistent, with $\hat{\rho}_x = 0.87$. Under the full-information permanent-income logic, households should therefore respond much more to aggregate income shocks than to idiosyncratic shocks.

They do not. The estimated pass-through of aggregate shocks is $\hat{\phi}_A = 0.126$, while the pass-through of idiosyncratic shocks is $\hat{\phi}_x = 0.118$. We cannot reject $H_0 : \phi_A = \phi_x$ (p-value $= 0.94$). In the data, households respond almost identically to a dollar of aggregate income shock and a dollar of idiosyncratic income shock, even though the aggregate component is far more persistent. This near-equality is difficult to reconcile with a standard full-information benchmark.

An obvious candidate explanation is household heterogeneity in wealth and liquidity. Hand-to-mouth households fully pass through income shocks regardless of their source or persistence. Hence, a higher mass of constrained households pushes both pass-through coefficients upward and compresses their ratio. But this force is quantitatively insufficient to explain the pattern we document. The mass of households that are so constrained must be

exceedingly large to push the average towards unity. That is hard to reconcile with the observed distribution of assets in the economy. Calibrating a full-information Bewley-Huggett economy to the estimated income process implies a much larger response to aggregate shocks than to idiosyncratic shocks: our calibrated model predicts $\phi_A/\phi_x \approx 4.4$, compared with a ratio close to one in the data. The puzzle, therefore, survives the standard incomplete-markets benchmark.

We argue that the missing ingredient is incomplete information about the source of income fluctuations. Households observe total income realizations, but they do not observe directly how much of a given realization reflects an aggregate component and how much reflects idiosyncratic variation or transitory noise. If households must infer the aggregate component from a noisy signal, they underreact to its future- income content. This mechanism attenuates the aggregate channel while leaving the idiosyncratic channel largely unchanged, compressing the ratio ϕ_A/ϕ_x toward one. This interpretation is supported by a growing literature showing that households face non-trivial inference problems when interpreting income movements and that subjective beliefs about earnings risk differ materially from realized-income estimates ((Carroll et al., 2020; Adams and Rojas, 2024; Caplin et al., 2023)).

We make this mechanism transparent in a simple two-period model and then discipline it quantitatively in a heterogeneous-agent economy. In the two-period CARA illustration, incomplete information replaces the aggregate present-value factor $(1 + \rho_A)$ with $(1 + \rho_A\bar{k})$, where $\bar{k} \in (0, 1)$ is a type-weighted Kalman gain, while leaving the idiosyncratic factor $(1 + \rho_x)$ unchanged. This delivers a clean intuition for why incomplete information compresses ϕ_A/ϕ_x. We then embed the same signal-extraction problem in a Bewley-Huggett economy with borrowing constraints and general equilibrium interest rate determination. The model is solved using the sequence-space Jacobian method of Auclert et al. (2021), augmented to accommodate information frictions. Using the empirically estimated income-process parameters and calibrating β to match the empirical idiosyncratic pass-through $\phi_x = 0.118$, the incomplete- information model produces $\phi_A/\phi_x \approx 1.09$, close to the empirical value and far below the full-information benchmark.

The mechanism also has important aggregate implications. Under full information, a one-standard-deviation aggregate productivity shock generates an impact consumption re-

sponse of roughly 12% of steady-state consumption in general equilibrium. Under incomplete information, the corresponding response falls to 2.8%. When households cannot cleanly distinguish persistent aggregate income movements from noise, they save rather than spend much of the shock, substantially dampening the aggregate demand response.

This paper contributes to three strands of the literature.

Pass-through of income shocks A large empirical literature estimates how consumption responds to permanent and transitory income shocks (Meghir and Pistaferri, 2004; Blundell et al., 2008; Kaplan and Violante, 2010; Heathcote et al., 2014; Kaplan and Violante, 2014; Arellano et al., 2017). Our contribution is to identify the aggregate and idiosyncratic persistent components of income separately and to test whether their relative pass-through aligns with the full-information PIH prediction. To our knowledge, we are the first to document and quantify the puzzle that $\phi_A \approx \phi_x$ despite $\rho_A \gg \rho_x$.

Income process estimation. Our estimation of the income process builds on the extensive literature that decomposes household earnings into persistent and transitory components (Meghir and Pistaferri, 2004, 2011; Guvenen, 2009; Guvenen and Smith, 2014). A key challenge in this literature is the identification of the persistent component and its volatility. Early work identified a permanent component by exploiting the unit root restriction; subsequent work relaxed this assumption and estimated a more general AR(1) structure. Our estimation goes further by decomposing the persistent component into an aggregate factor, with heterogeneous household exposures λ_i, and an idiosyncratic component. This separation is conceptually motivated by Guvenen et al. (2014), who show that exposure to aggregate cycles is systematically different for people of different demographics. Our dynamic factor approach recovers both the aggregate income factor and the household-specific loadings from the cross-sectional covariance structure of income, avoiding the confound between aggregate and idiosyncratic persistence that plagues single-component models. Our estimation closely follows Braxton et al. (2025) by combining a Kalman filter with the EM algorithm to recover latent income states from panel data. Relative to their generalized filter, our contribution is to adapt this approach to a setting with an explicit aggregate factor structure, so that household income is decomposed into an aggregate persistent component and an idiosyncratic persistent component rather than only persistent and temporary individual earnings risk.

Incomplete information in macroeconomics. Our paper closely relates to the literature on informational frictions and aggregate demand (Lorenzoni, 2009; Angeletos and La'O, 2013; Angeletos and Lian, 2022). The idea that households underreact to aggregate income shocks owing to a lack of information dates back to Pischke (1995). This paper proposes this friction as an explanation for excess smoothness in aggregate consumption series. We formalize this argument with empirical backing and argue that information frictions are indeed important to account for in quantitative macro models studying aggregate demand responses to aggregate income shocks. Guvenen (2007) highlights that the distinction between a highly persistent (but stationary) component and a random walk is particularly consequential: households who are uncertain about this distinction will behave differently from those under full information, since their estimate of the permanent income gain from a given shock depends critically on what they believe about its persistence. In this paper, households in our economy are fully aware of the parameters of the income process but fail to distinguish the sources of shocks. More recently, Adams and Rojas (2024) studied a similar inference problem in a small, open-economy, heterogeneous-agent model, assuming an income process in which idiosyncratic shocks are more persistent than aggregate shocks. In our closed-economy Huggett framework, the estimated persistence ranking runs in the opposite direction, so incomplete information attenuates rather than amplifies aggregate consumption responses; moreover, both the persistence ordering and the signal-to-noise ratio are pinned down by the estimated household-income process for US households.

Heterogeneous-agent models and aggregate demand. Our quantitative analysis builds on the literature studying how household heterogeneity shapes aggregate consumption dynamics (Kaplan et al., 2018; Auclert, 2019; McKay and Reis, 2016; Guerrieri and Lorenzoni, 2017). These papers emphasize the role of *who* receives the income; high-MPC households, such as hand-to-mouth consumers, amplify fiscal and monetary transmission. Our paper adds a complementary dimension: not only does who receives the income matter, but what households can infer from it. Even among unconstrained households with standard intertemporal motives, the information environment determines how much of the aggregate income signal is acted upon, with first-order implications for aggregate demand.

The remainder of the paper is organized as follows. Section 2 establishes the empirical facts on persistence and pass-through. Section 3 presents the full-information quantitative benchmark and shows that standard household heterogeneity does not resolve the puzzle. Section 4 introduces a simple two-period model that isolates the information-friction mechanism. Section 5 embeds that mechanism in the quantitative Bewley–Huggett economy and highlights the implications for aggregate demand. Section 6 concludes.

2 Empirical Evidence

How does household consumption respond to aggregate and idiosyncratic income shocks that differ in persistence? To answer this question, we use data from the Panel Study of Income Dynamics (PSID) to estimate an income process with persistent aggregate and idiosyncratic shocks and their respective pass-through to household consumption. The consumption pass-through coefficient measures the extent to which an income innovation translates into a change in household consumption. We find that although aggregate income shocks are more persistent than idiosyncratic shocks, their estimated consumption pass-through coefficients are similar. This finding remains robust when we exclude liquidity-constrained households, suggesting that the similarity in consumption responses cannot be explained solely by households' limited ability to smooth consumption.

2.1 Estimation of Income Shocks

To understand how households respond to different types of income shocks, we estimate an income process that incorporates both aggregate and idiosyncratic shocks with different degrees of persistence. We use household income and consumption data from eleven biennial waves of the Panel Study of Income Dynamics (PSID), spanning 1999 to 2019. We restrict the sample to households whose heads are between 30 and 65 years of age, exclude households with annual income below $5,000, and require at least five observed income

waves per household.[1] The resulting sample consists of 1,908 households and 14,819 observations.

We residualize household income and consumption by regressing log labor income and log consumption on a rich set of observables, including a third-order polynomial in age and indicators for race, region, race-year, region-year, family size, year of birth, number of children, and outside dependents. We interpret the resulting residuals as the unexplained components in household income and consumption. The variances of residualized income and consumption are approximately 0.26 and 0.12, respectively, consistent with estimates in the literature.[2] The covariance between income and consumption residuals is 0.004 and statistically significant, providing an important moment for identifying the pass-through of income shocks to consumption.[3]

We model residualized log labor income as the sum of a persistent aggregate component, a persistent idiosyncratic component, and a transitory component. Specifically, residualized income, z_{it}, is given by

$$z_{it} = \lambda_i A_t + x_{it} + m_{it}, \quad m_{it} \sim \mathcal{N}(0, \sigma_m^2), \tag{1}$$

where A_t is a common aggregate factor with heterogeneous household exposure λ_i, x_{it} is a persistent idiosyncratic component, and m_{it} captures transitory income shocks. The aggregate and idiosyncratic components follow AR(1) processes:

$$\begin{aligned} A_t &= \rho_A A_{t-1} + \epsilon_t^A, \quad \epsilon_t^A \sim \mathcal{N}(0, \sigma_A^2), \\ x_{it} &= \rho_x x_{it-1} + \epsilon_{it}^x, \quad \epsilon_{it}^x \sim \mathcal{N}(0, \sigma_x^2). \end{aligned} \tag{2}$$

The parameters ρ_A and ρ_x govern the persistence of aggregate and idiosyncratic income shocks, respectively, while σ_A^2 and σ_x^2 denote their shock variances. We allow households to

[1] We provide additional details on the data and sample selection in Appendix A. Table A1 shows summary statistics of the estimation sample.

[2] Table A2 presents the empirical moments calculated using residualized income and consumption. We refer the reader to Ghosh and Theloudis (2023) for a detailed comparison of these moments with the literature.

[3] Both income and consumption residuals exhibit high kurtosis, consistent with substantial tail risk. Ghosh and Theloudis (2023) exploit these higher-order moments to discipline consumption pass-through coefficients under a nonlinear pass-through specification.

differ in their exposure to aggregate shocks through λ_i, while assuming that the persistence and innovation variance of idiosyncratic shocks, ρ_x and σ_x^2, and the variance of transitory shocks, σ_m^2, are common across households.

We estimate the income process in two stages to separately identify the aggregate and idiosyncratic components of household income. In the first stage, we exploit the common variation in household income to estimate the aggregate factor and its persistence, ρ_A. In the second stage, we remove the estimated aggregate component from household income and use the remaining idiosyncratic variation to estimate the persistence of idiosyncratic shocks, ρ_x, and the shock variances, σ_x^2 and σ_m^2.

First Stage: We estimate the aggregate component using a factor structure:

$$\begin{aligned} z_{it} &= \lambda_i A_t + s_{it} \\ A_t &= \rho_A A_{t-1} + \epsilon_t^A \end{aligned} \tag{3}$$

where A_t captures common fluctuations in household income, λ_i measures household-specific exposure to these fluctuations, and s_{it} captures the remaining idiosyncratic variation. The key identification assumption is that aggregate shocks generate common movements in income across households, whereas idiosyncratic shocks are household-specific. We exploit this common variation to extract the first principal component of residualized income as an estimate of A_t. The factor model also yields household-specific loadings, λ_i, which determine the contribution of the aggregate factor to each household's income. We then fit an AR(1) process to the estimated aggregate factor to recover its persistence, ρ_A, and the corresponding innovations, ϵ_t^A.

An important identification consideration is that the aggregate factor and its loadings are not separately identified without a normalization. Specifically, the factor model identifies the product $\lambda_i A_t$, but not the scale of A_t independently of λ_i. Consequently, while the persistence parameter ρ_A is identified, the aggregate shock variance σ_A^2 depends on the normalization of the factor. We therefore focus on the persistence of the aggregate component, which is sufficient for our analysis. Despite the relatively short time dimension of the panel, we find that the estimated persistence is stable across bootstrap replications. Panel A of Table

Table 1: Income Process Parameters

	Estimates	
	With Aggregate Factor	No Aggregate Factor
	(1)	(2)
A. Persistence Parameters		
ρ_A	0.996	–
	(0.003)	–
ρ_x	0.871	0.955
	(0.023)	(0.004)
B. Variance parameters		
σ_x^2	0.058	0.041
	(0.006)	(0.003)
σ_m^2	0.046	0.058
	(0.004)	(0.003)

Notes: This table presents the parameters from the income process estimation. Panel A shows the persistence of the aggregate and idiosyncratic components. Panel B shows the variances estimated for each component. Column (1) shows the estimates of the income process with aggregate factor, and column (2) shows the estimates with just idiosyncratic shocks. Block bootstrap standard errors are reported in the parentheses.

1 reports an estimated aggregate persistence of 0.99, indicating that aggregate income shocks are highly persistent, consistent with evidence in the literature on persistent and permanent income components.

Second Stage: Having estimated the aggregate component, we subtract its contribution from household income to isolate the idiosyncratic variation, s_{it}. We model this remaining variation as the sum of a persistent idiosyncratic component, x_{it}, and a transitory component, m_{it}. We estimate these components using the following state-space representation:

$$\begin{aligned} \text{Measurement:} \quad & s_{it} = x_{it} + m_{it}, \\ \text{State:} \quad & x_{it} = \rho_x x_{it-1} + \epsilon_{it}^x. \end{aligned} \tag{4}$$

We estimate the parameters $\rho_x, \sigma_x^2, \sigma_m^2$ and the latent states x_{it} using a Kalman filter combined with the Expectation-Maximization (EM) algorithm, following Braxton et al. (2025) closely. The algorithm iterates between two steps: the E-step computes the expected sufficient statis-

tics of the latent states using Kalman smoothing, while the M-step updates the model parameters by maximizing the expected complete-data log-likelihood. Iterating between these steps yields estimates of the persistence and variances of the idiosyncratic income components, along with estimates of the underlying idiosyncratic innovations, ϵ_{it}^{x}.[4] We provide additional details on the estimation procedure in Appendix A.

Panel A of Table 1 reports an estimated annual persistence of 0.99 for the aggregate component, ρ_A, compared with 0.87 for the idiosyncratic component, ρ_x, indicating that aggregate income shocks are more persistent than idiosyncratic shocks. Although we estimate the income process using biennial data, we annualize the persistence parameters to facilitate comparison with the literature.[5] When we exclude the aggregate factor from the income process, the estimated idiosyncratic persistence increases from 0.87 to 0.95, consistent with estimates commonly reported in the consumption literature. This comparison suggests that omitting aggregate income fluctuations leads to an overestimation of idiosyncratic income persistence, as persistent aggregate fluctuations are partially attributed to the idiosyncratic component. Panel B reports the estimated shock variances. When accounting for the aggregate factor, the variances of persistent idiosyncratic and transitory shocks are 0.058 and 0.046, respectively. Excluding the aggregate factor reduces the estimated variance of persistent idiosyncratic shocks to 0.041 while increasing the variance of transitory shocks to 0.058. These estimates are broadly consistent with those reported in the existing literature.

To formally assess whether aggregate income shocks are more persistent than idiosyncratic shocks, we test the null hypothesis $H_0 : \rho_A \leq \rho_x$ against the alternative $H_1 : \rho_A > \rho_x$. The estimated difference in persistence is $\rho_A - \rho_x = 0.125$, with a standard error of 0.024. We compute the standard error using 1,000 paired household-bootstrap replications, re-estimating both persistence parameters in each replication to account for the covariance between their estimates. The resulting p-value is below 0.001, strongly rejecting the null hypothesis. This result provides statistically significant evidence that aggregate income shocks are more persistent than idiosyncratic shocks.

[4] See Braxton et al. (2025) for details on the estimation procedure and identification of parameters.

[5] We annualize the estimated biennial persistence parameters using $\rho_{\text{annual}} = \sqrt{\rho_{\text{biennial}}}$, consistent with an underlying annual AR(1) process.

2.2 Consumption Pass-through Coefficients

Having recovered the household-level income shocks from the estimated income process, we next examine how these shocks transmit to household consumption. Let c_{it} denote residualized log consumption, after removing the variation explained by household observables as described above. We define the consumption pass-through coefficients as the covariance between consumption growth and income shocks, normalized by the variance of the corresponding income shocks:

$$\begin{aligned} \phi_A &= \frac{\mathrm{Cov}(\Delta c_{it}, \lambda_i \epsilon_t^A)}{\mathrm{Var}(\lambda_i \epsilon_t^A)}, \\ \phi_x &= \frac{\mathrm{Cov}(\Delta c_{it}, \epsilon_{it}^x)}{\mathrm{Var}(\epsilon_{it}^x)}. \end{aligned} \tag{5}$$

Here, ϕ_A and ϕ_x measure the consumption pass-through of aggregate and idiosyncratic persistent income shocks, respectively. We scale aggregate innovations by the household-specific factor loading, λ_i, to account for heterogeneity in households' exposure to aggregate shocks. This allows us to measure both types of shocks in terms of their contributions to household income.

We estimate the consumption pass-through coefficients using the following regressions:

$$\begin{aligned} \Delta c_{it} &= \mu_i^A + \delta_t^A + \phi_A \widehat{(\lambda_i \epsilon_t^A)} + error \\ \Delta c_{it} &= \mu_i^x + \delta_t^x + \phi_x \widehat{\epsilon_{it}^x} + error \end{aligned} \tag{6}$$

where μ_i^j and δ_t^j denote household and time fixed effects, respectively, for $j \in A, x$. The estimated coefficients measure the sensitivity of household consumption growth to innovations in each income component. In particular, a pass-through coefficient of ϕ_j implies that a one-percent innovation in household income arising from component j is associated with an approximately ϕ_j percent change in consumption. These coefficients also provide a measure of consumption insurance against income shocks: a lower pass-through indicates that households absorb a smaller fraction of income fluctuations through adjustments in consumption, consistent with greater consumption smoothing.

Table 2 reports the estimated consumption pass-through coefficients. On average, households pass through approximately 13% of aggregate income shocks and 12% of idiosyncratic

Table 2: Estimates of Pass-through Coefficients

	Pass-through Estimates	
	With Aggregate Factor	No Aggregate Factor
	(1)	(2)
ϕ_A: Aggregate Pass-through	0.13	-
	(0.08)	-
ϕ_x: Idiosyncratic Pass-through	0.12	0.18
	(0.03)	(0.02)
	Test Statistics	
	With Aggregate Factor	No Aggregate Factor
p-value of $H_0 : \phi_A = \phi_x$	0.94	-
Observations	12,543	12,543
Households	1,907	1,907

Notes: This table presents the coefficients estimated by running a regression of consumption change on estimated income shocks. Column (1) shows the estimates when the aggregate factor is included in the estimation procedure, whereas column (2) shows the estimates without the aggregate factor. Block bootstrap standard errors are reported in the parentheses.

income shocks to consumption. These estimates indicate substantial consumption smoothing against both types of shocks, with nearly identical responses despite the significantly higher persistence of aggregate shocks. However, when we exclude the aggregate factor from the income process, the estimated pass-through of idiosyncratic shocks increases from 0.12 to 0.18. The estimated pass-through coefficients are consistent with those reported in Ghosh and Theloudis (2023), which range from 0.11 to 0.16. Together with the income process estimates, this finding highlights the importance of accounting for heterogeneous exposures to aggregate income fluctuations when estimating the pass-through of income shocks to household consumption.

2.3 Robustness and Discussion

The similarity in the consumption pass-throughs to aggregate and idiosyncratic income shocks presents a puzzle for the canonical full-information benchmark. In the next section,

we show that a standard incomplete-markets Huggett model predicts larger consumption responses to more persistent income shocks. Given the significantly higher persistence of aggregate income shocks, the benchmark predicts that $\phi_A > \phi_x$. However, the estimated pass-through coefficients of 0.13 and 0.12 for aggregate and idiosyncratic shocks, respectively, differ only marginally. To formally assess whether these coefficients differ, we test $H_0 : \phi_A = \phi_x$ against $H_1 : \phi_A \neq \phi_x$. The resulting p-value of 0.94 indicates that we cannot reject equality of the two coefficients at conventional significance levels. This finding motivates our investigation into why consumption responses exhibit such limited sensitivity to differences in income persistence.

We further examine whether the similarity in consumption pass-through is driven by liquidity-constrained households, whose limited ability to smooth income fluctuations may lead to similar consumption responses across shocks of different persistence. To assess this possibility, we first exclude households in the bottom income quartile and then exclude those in the bottom wealth decile, since these households are more likely to face liquidity constraints. Panel A of Table B1 in Appendix B reports the results. Excluding low-income households yields pass-through coefficients of 0.104 and 0.122 for aggregate and idiosyncratic shocks, respectively. Further excluding low-wealth households produces estimates of 0.088 and 0.116. Thus, even among households less likely to face liquidity constraints, the consumption response to idiosyncratic shocks slightly exceeds that to aggregate shocks. This suggests that the absence of a substantially larger response to more persistent aggregate shocks cannot be explained solely by liquidity constraints.

We also examine the robustness of our results to the length of household income histories and the presence of outliers. While our baseline sample requires at least five observed income waves per household, we consider more restrictive samples requiring at least seven and nine waves. We additionally winsorize the data at the 1% level to assess the sensitivity of our estimates to extreme observations. Panel B of Table B1 reports the results. Across these specifications, the estimated pass-through coefficients remain broadly similar for aggregate and idiosyncratic shocks, reinforcing our main finding that differences in income persistence are not reflected in comparably large differences in consumption responses.

2.4 Consumption Impulse Responses to Income Shocks

The consumption pass-through coefficients estimated above capture the contemporaneous response of household consumption to income innovations. To examine how these shocks are associated with consumption growth over subsequent periods, we estimate dynamic responses using the local projection approach of Jordà (2005) and Jordà and Taylor (2025). This approach allows us to estimate the consumption response directly at each horizon, without imposing a particular functional form on its dynamics.

For each horizon $h \in 0, 1, \ldots, 5$, we estimate the following regression:

$$\Delta c_{i,t+h} = \mu^k_{i,h} + \delta^k_{t,h} + \beta_{k,h}\widehat{g}^k_{it} + u^k_{i,t,h}, \quad \text{for } k \in A, x \tag{7}$$

where $\mu^k_{i,h}$ and $\delta^k_{t,h}$ denote household and time fixed effects, respectively, and $\widehat{g}it^k$ represents the estimated income shock. Specifically, the aggregate and idiosyncratic income innovations are constructed as

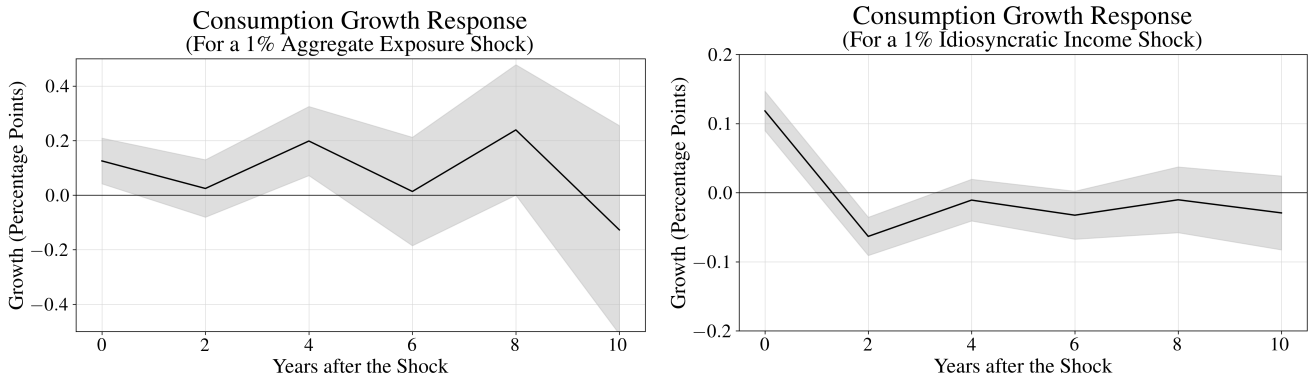
$$\widehat{g}^A_{it} = \widehat{\lambda}_i\big(\widehat{A}_t - \widehat{\rho}_A\widehat{A}_{t-1}\big), \qquad \widehat{g}^x_{it} = \widehat{x}_{it|T} - \widehat{\rho}_x\widehat{x}_{i,t-1|T}.$$

where $\widehat{x}it|T$ denotes the smoothed estimate of the persistent idiosyncratic income component, and the persistence parameters used to construct the innovations correspond to the biennial frequency of the data. We estimate separate regressions for each type of income shock and horizon, without including lagged consumption growth or additional controls. At the impact horizon, $h = 0$, the estimated coefficients coincide with the contemporaneous pass-through coefficients, $\widehat{\beta}_{A,0} = \widehat{\phi}_A$ and $\widehat{\beta}_{x,0} = \widehat{\phi}_x$.

Since the PSID data are biennial, the estimated horizons correspond to 0, 2, 4, 6, 8, and 10 years after the shock. We normalize the responses to a positive 1% innovation in household income, so that $\widehat{\beta}_{k,h}$ measures the associated percentage-point change in biennial consumption growth at horizon h.[6]

[6]The local projections estimate dynamic associations between consumption growth and recovered income innovations. A causal interpretation additionally requires appropriate exogeneity assumptions. In particular, the smoothed idiosyncratic innovations use the full income history, and the separate regressions do not control for the other contemporaneous income innovation.

Figure 1: Impulse Response to Income Shocks



(a) Aggregate Exposure Shocks

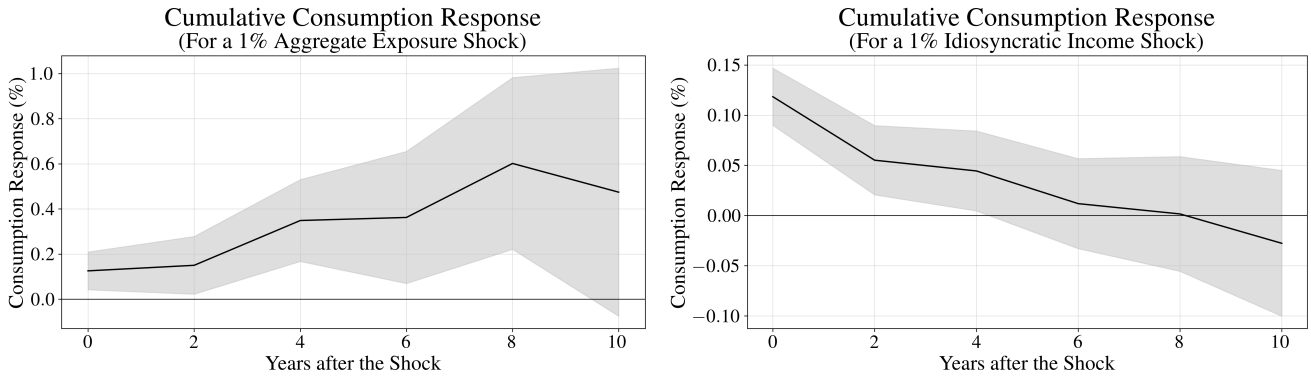
(b) Idiosyncratic Shocks

Notes: The figure reports local projection estimates of the response of residualized biennial consumption growth to a positive 1% household income innovation. Panel (a) presents responses to aggregate income innovations, scaled by household-specific factor loadings, while Panel (b) presents responses to persistent idiosyncratic income innovations. All regressions include household and wave fixed effects and are estimated separately at horizons of 0, 2, 4, 6, 8, and 10 years. Shaded areas represent pointwise 66% normal-approximation confidence bands constructed using standard errors from 1,000 full-estimation household-bootstrap replications.

Figure 1 presents the dynamic consumption responses to aggregate and idiosyncratic income shocks. Although the impact responses are nearly identical, the subsequent adjustment patterns differ substantially. Following an idiosyncratic income shock, consumption growth increases on impact but turns negative in the next biennial period, suggesting that households partially reverse their initial consumption response. This pattern is consistent with an initial overreaction to idiosyncratic income shocks, followed by a subsequent correction. In contrast, consumption responds positively to aggregate income shocks on impact and exhibits additional positive growth responses at several subsequent horizons, suggesting a more gradual adjustment to aggregate income fluctuations. This pattern is consistent with an initial underreaction to aggregate shocks, with households adjusting their consumption over time as the effects of these shocks unfold. Although the aggregate responses are estimated with considerable uncertainty at longer horizons, these patterns provide suggestive evidence of differences in the timing of consumption adjustments across the two types of income shocks.

Figure 2 presents the cumulative consumption responses to aggregate and idiosyncratic income shocks, constructed by summing the estimated consumption growth responses across horizons. Although the contemporaneous pass-through coefficients are nearly identical, the

Figure 2: Cumulative Impulse Response to Income Shocks



(a) Aggregate Exposure Shocks

(b) Idiosyncratic Shocks

Notes: The figure presents cumulative consumption responses to a positive 1% household income innovation, estimated using local projections following Jordà (2005). Panel (a) reports responses to aggregate income shocks, scaled by household-specific factor loadings, while Panel (b) reports responses to persistent idiosyncratic income shocks. Cumulative responses are constructed by summing the estimated consumption growth responses from the impact period through each horizon. All regressions include household and wave fixed effects and are estimated separately at biennial horizons spanning 0 to 10 years after the shock. Responses are expressed as percentage changes in consumption relative to the pre-shock level. Shaded areas represent pointwise 66% confidence bands based on 1,000 full-estimation household-bootstrap replications, accounting for covariance across horizons in the cumulative responses.

cumulative responses reveal strikingly different adjustment patterns. Following an idiosyncratic income shock, consumption initially increases by approximately 0.12%, but this response gradually reverses, with the cumulative effect approaching zero after eight years. In contrast, the cumulative response to an aggregate income shock continues to increase beyond the impact period, reaching over 0.50% after eight years, compared with an initial response of 0.13%. This result strengthens our claim that households initially overreact to idiosyncratic income shocks and subsequently correct their consumption, whereas they exhibit a more gradual adjustment to aggregate shocks. While the longer-horizon aggregate responses are estimated with considerable uncertainty, the contrasting dynamics provide suggestive evidence that similar contemporaneous pass-through coefficients conceal important differences in how households adjust consumption over time.

These empirical findings motivate our quantitative analysis of consumption responses in a full-information benchmark. In the next section, we calibrate a standard incomplete-markets Huggett economy and examine its general-equilibrium implications for consumption pass-through to aggregate and idiosyncratic income shocks. We show that the model

predicts a substantially larger consumption response to the more persistent aggregate shocks, generating a much greater difference in pass-through coefficients than observed in the data. This discrepancy highlights the difficulty of reconciling the empirical evidence with the canonical full-information framework.

3 Quantitative Full-Information Benchmark

In this section, we examine whether a standard incomplete-markets model, disciplined by the income-process estimates from Section 2, can account for the nearly identical consumption pass-through coefficients for aggregate and idiosyncratic income shocks observed in the data. We show that the model fails to reproduce this empirical finding. We first present a quantitative Bewley–Huggett economy under full information and then describe the sequence-space Jacobian approach used to compute consumption pass-through coefficients and general-equilibrium responses. Finally, we show that the benchmark predicts substantially greater consumption pass-through for aggregate shocks than for idiosyncratic shocks, $\phi_A \gg \phi_x$, in sharp contrast to the empirical evidence.

3.1 Model Setup

We consider an incomplete-markets economy in which households choose consumption, saving, and labor supply. Households face persistent and transitory idiosyncratic productivity risk and differ in their exposure to aggregate productivity. In the model, households also face a borrowing constraint. Firms produce with effective labor under perfect competition, so profits are zero. The government taxes labor earnings and rebates all revenue through a lump-sum transfer. Endogenous hours and redistribution act as insurance mechanisms for the households facing income shocks.

Household Preferences and Budget Constraint. Time is annual and a unit measure of infinitely-lived households indexed by $i \in [0, 1]$ have time-separable preferences over con-

sumption c_{it} and hours worked n_{it}:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(c_{it}, n_{it}), \qquad \beta \in (0,1). \tag{8}$$

The expectation $\mathbb{E}_0$ is conditional on the household's information set at time zero, which depends on the information environment. We assume the preferences are

$$u(c_{it}, n_{it}) = \left[u(c_{it}) - \psi \frac{n_{it}^{1+\varphi}}{1+\varphi} \right], \qquad \psi, \varphi > 0, \tag{9}$$

where $c_{it} > 0$, hours satisfy $0 \leq n_{it} \leq 1$, and

$$u(c) = \begin{cases} \dfrac{(c^{1-\gamma} - 1)}{(1-\gamma)}, & \gamma \neq 1 \text{ and } \gamma > 0, \\ \log c, & \gamma = 1. \end{cases} \tag{10}$$

The intertemporal elasticity of substitution (IES) is $1/\gamma$, and the Frisch elasticity is $1/\varphi$. Hours are a fraction of a unit time endowment. Households hold a single liquid bond b_{it} paying the gross return $1 + r_t$, which they take as given. Households pay taxes on their labor income at a rate τ and receive a lump-sum transfer from the government every period T. The period budget constraint is

$$c_{it} + b_{i,t+1} = (1 + r_t)\, b_{it} + (1 - \tau) y_{it} + T_t, \tag{11}$$

subject to the borrowing constraint $b_{i,t+1} \geq \underline{b}$. Labor income, y_{it}, is realized and observed each period.

Labor income. Let w_t denote the competitive wage per efficiency unit. The gross labor earnings also depend on the households labor efficiency $\exp(z_{it})$:

$$y_{it} = w_t n_{it} \exp(z_{it}) \tag{12}$$

We assume the income process to be similar to our empirical part. We assume the household labor efficiency follows

$$z_{it} = \lambda_i A_t + x_{it} + m_{it} \tag{13}$$

The fixed loading λ_i captures direct exposure to aggregate productivity. We represent exposure using fixed loading types, $\{\lambda_1, \ldots, \lambda_J\}$, with population shares $\{\omega_1, \ldots, \omega_J\}$. The persistent states follow

$$\begin{aligned} A_t &= \rho_A A_{t-1} + \epsilon_t^A, \quad \epsilon_t^A \sim \mathcal{N}(0, \sigma_A^2), \\ x_{it} &= \rho_x x_{i,t-1} + \epsilon_{it}^x, \quad \epsilon_{it}^x \sim \mathcal{N}(0, \sigma_x^2), \end{aligned} \tag{14}$$

where $0 < \rho_A, \rho_x < 1$. The transitory state satisfies $m_{it} \sim \mathcal{N}(0, \sigma_m^2)$ and is independent over time. Shocks are mutually independent, and idiosyncratic shocks are independent across households. Idiosyncratic risk parameters are common across loading types. Finally, the household i's log labor income can be written as

$$\log y_{it} = \log w_t + \log n_{it} + \lambda_i A_t + x_{it} + m_{it}, \tag{15}$$

Production. A representative competitive firm produces the final consumption good using aggregate effective labor as the only input. The production technology is

$$Y_t = \exp(A_t) N_t, \qquad N_t = \int \exp(z_{it}) n_{it} \, di. \tag{16}$$

where $\exp(A_t)$ is aggregate total factor productivity, increasing in the aggregate earnings state A_t, and N_t is aggregate effective labor input.

Taking A_t and the wage w_t as given, the firm maximizes profits. The first-order condition for labor demand pins down the competitive wage:

$$w_t = \exp(A_t). \tag{17}$$

Given the factor loading λ_i in (15), the total exposure of an aggregate shock to household i's hourly compensation is $1 + \lambda_i$.

Government. The government has no debt or purchases and rebates all labor-tax revenue equally. Because the population has unit mass, its budget is

$$T_t = \tau \int y_{it} di = \tau w_t N_t = \tau Y_t. \tag{18}$$

Note that the full rebate condition determines the lump-sum transfers endogenously.

Market Clearing. There is a single asset, the bond, which is in zero net supply. Bond-market clearing requires

$$\int b_{it+1}\, di = 0 \tag{19}$$

in every period, and the equilibrium interest rate r_{t+1} adjusts to satisfy this condition. Since there is no investment or government spending in the model, the goods market clears when aggregate output equals aggregate consumption:

$$Y_t = C_t \equiv \int c_{it}\, di. \tag{20}$$

3.2 Equilibrium under Full Information

In this full information benchmark, households observe current productivity states, aggregate productivity, prices and transfers. Households take prices, the transfers and the aggregate law of motion as given. The complete individual state vector is $(b_{it}, x_{it}, m_{it}, \lambda_i, \Gamma_t)$, where Γ_t is the cross-sectional distribution over individual states $(b_i,\ x_i, m_i, \lambda_i)$. The equilibrium interest rate r_{t+1} is determined by bond-market clearing, which depends on aggregate bond demand. The aggregate state is $\mathcal{S}_t = (A_t, r_t, \Gamma_t)$, with lump-sum transfers $T_t = T(\mathcal{S}_t)$. The household solves

$$\begin{aligned} V(b, x, m, \lambda; \mathcal{S}) = \max_{c>0,\, 0\le n\le 1,\, b'\ge \underline{b}} \Big\{ & u(c) - \psi \frac{n^{1+\varphi}}{1+\varphi} \\ & + \beta \mathbb{E}\big[V(b', x', m', \lambda; \mathcal{S}') \mid x, \mathcal{S}\big] \Big\}, \end{aligned} \tag{21}$$

subject to (11).

A recursive competitive equilibrium consists of household value functions, consump-

tion, saving, and labor policy functions, the transfer rule, and consistency of the distribution's transition law, a set of prices w_t and r_t such that: (i) households optimize, solving (21), given $\{r_t, w_t\}$ and their information set; (ii) the firm optimizes, with wages satisfying (17); (iii) government budget balances (18); (iv) bond-market clearing (19) and goods-market clearing (20) both hold.

3.3 Calibration

To calibrate the model, we start at the stationary equilibrium with no aggregate shocks ($A_t = 0$) but keep idiosyncratic shocks. Given the aggregate state variables, we solve household policies and the invariant distribution. We adjust the return until mean bonds are zero and the transfer until it equals tax revenue. The household solver uses the endogenous-grid method of Carroll (2006).[7] Table 3 distinguishes fixed parameters from calibrated parameters. Following Kaplan et al. (2018), we set both IES and Frisch elasticities to one and fix $\tau = 0.30$ for the labor-tax rate. The borrowing constraint is set too loose to allow for lower pass-through coefficients observed in the data. We use the empirical estimates as direct inputs to the parameters governing idiosyncratic shocks, $\{\rho_x, \sigma_x, \sigma_m\}$.

We calibrate the parameters (β, ψ) to match the average share of hours at 0.5, following Kaplan et al. (2018), and the pass-through coefficient for idiosyncratic shocks, based on our empirical evidence. This is a joint calibration exercise, not a separate one-to-one mapping between each preference and moment. The model solution gives the equilibrium quantities for the real interest rate and transfers.

$$\mathbb{E}_\Gamma[n] = 0.5, \quad \phi_x^{\text{model}} = 0.118 \tag{22}$$

To compute the pass-through statistic for idiosyncratic shocks, we observe complete model histories at two-year intervals, fit an auxiliary stationary AR(1)-plus-transitory process to log earnings, smooth its persistent state and construct its innovations. We regress two-year log consumption growth on these generated earnings innovations with household and

[7] We provide more details on the solution method in Appendix C. The quantitative baseline imposes feasibility on this finite productivity support: even the lowest wage state permits positive consumption at the borrowing limit. The fixed borrowing limit does not assert feasibility at arbitrarily low wages.

Table 3: Baseline stationary calibration

Description	Value	Source/Reference
A. Fixed inputs		
Intertemporal elasticity, $1/\gamma$	1.000	Kaplan et al. (2018)
Frisch elasticity, $1/\varphi$	1.000	Kaplan et al. (2018)
Labor-income tax rate, τ	0.300	Income tax rate
Borrowing limit, $\underline{b}$	-10.000	Constrained-HH
Annual persistent productivity, ρ_x	0.871	Estimate from income process
Annual persistent innovation SD, σ_x	0.181	Estimate from income process
Transitory productivity SD, σ_m	0.215	Estimate from income process
B. Calibrated parameters		
Annual discount factor, β	0.976	Average hours and pass-through
Labor-disutility weight, ψ	2.515	Average hours and pass-through

Notes: This table shows the parameters that are set exogenously and calibrated in the model. Both preference parameters jointly match mean hours of 0.5 and biennial earnings-based pass-through of 0.118. The return and transfer jointly satisfy bond market and government budget clearing conditions.

time effects. During calibration, we validate the fit with two independent panels of 20,000 households and 201 annual dates each, initialized from the stationary distribution. Each panel contains 100 two-year growth observations per household, thereby avoiding empirical missingness and short household histories. The consumption-growth interval, generated-earnings interpretation, and fixed effects retain the empirical definition of the response.

Panel B of Table 3 shows the calibrated parameters. The annual discount factor β is 0.976, and the labor disutility weight ψ is 2.515. Table 4 presents model outcomes along with the targeted moments. Mean hours and the population earnings response match their targets, and the independent simulated mean response is 0.118. The equilibrium net return is 2.25% annually. The borrowing limit binds for 0.37% of households, while 1.63% choose the hours ceiling.

Table 4: Calibration Targets and Model Outcomes

	Target	Model
A. Targeted Moments		
Mean hours, $\mathbb{E}_\Gamma[n]$	0.500	0.500
Biennial idiosyncratic pass-through, ϕ_x	0.118	0.118
B. Untargeted Model Outcomes		
Annual real interest rate, r		2.25%
Transfer, T		0.186
Mean gross annual earnings		0.620
Borrowing-limit share		0.37%
Hours-ceiling share		1.63%

Notes: This table shows the calibration targets and equilibrium model outcomes. Panel A presents the targeted moments used to calibrate β and ψ. Panel B shows equilibrium objects in the model.

3.4 Response to Aggregate Shocks

We set annual aggregate persistence to $\rho_A = 0.996$, using the annualized empirical estimate, and the annual innovation standard deviation to $\sigma_A = 0.002$. Table 5 presents the aggregate inputs and the distributional characteristics of exposures. We approximate a normal exposure distribution with 20 equally weighted types. Let $q_g = \Phi^{-1}((g-1/2)/20)$, where Φ is the standard normal distribution function. The loadings are

$$\lambda_g = 0.2 \frac{q_g - \bar{q}}{\left[20^{-1} \sum_{j=1}^{20} (q_j - \bar{q})^2\right]^{1/2}}, \qquad \omega_g = 1/20, \quad g = 1, \ldots, 20. \tag{23}$$

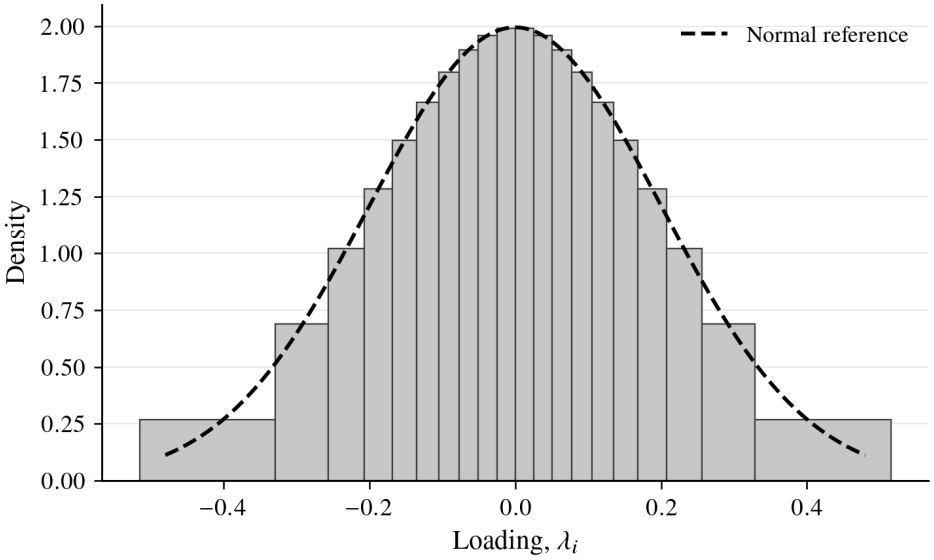
Thus the discrete distribution has mean zero and standard deviation 0.2 exactly. We assign exposure independently of initial assets, persistent productivity, transitory productivity, and income. Exposure is a permanent household characteristic. Each type initially has the same conditional household distribution; subsequent aggregate shocks can generate dependence between exposure, assets, and income. Figure 3 shows the distribution of the loadings, ranging from -0.405 to 0.405. Total hourly-compensation exposure, $1+\lambda_g$, is positive for all types,

Table 5: Aggregate shocks and household exposure

Parameter or statistic	Value	Assignment
Annual aggregate persistence, ρ_A	0.996	Empirical estimate
Annual innovation SD, σ_A	0.002	To match aggregate volatility
Mean direct exposure, $\mathbb{E}[\lambda_i]$	0	Normalization
Exposure SD	0.200	Estimation from income process
Number of exposure types	20	Equal-mass normal quantiles
Minimum total exposure,	0.595	Implied
Maximum total exposure,	1.405	Implied

Exposure types follow equation (23) and are assigned independently of initial household states. Each type has population mass 0.05, and its loading remains fixed. Neither aggregate shock variance nor exposure dispersion is a calibration target.

Figure 3: Distribution of fixed household exposure



Notes: The model uses 20 equally weighted discrete types, with mean zero and standard deviation 0.2. Histogram heights divide type mass by bin width; the dashed line is a normal reference density. Types are initially independent of assets, productivity, and income, and remain fixed as household states evolve.

ranging from 0.595 to 1.405. A positive innovation raises every type's hourly compensation on impact, holding its idiosyncratic state fixed. Endogenous hours, asset holdings, and equilibrium prices determine the earnings and consumption responses.

Solution method. At each date, the bond return and the transfers clear the bond market and government budget. For paths $\mathbf{R}$, $\mathbf{T}$ and $\mathbf{A}$, define

$$\mathcal{H}_t(\mathbf{R}, \mathbf{T}, \mathbf{A}) = \begin{pmatrix} B_{t+1}(\mathbf{R}, \mathbf{T}, \mathbf{A}) \\ T_t - \tau Y_t(\mathbf{R}, \mathbf{T}, \mathbf{A}) \end{pmatrix} = 0. \tag{24}$$

We solve a first-order sequence-space approximation around the calibrated stationary economy, following Auclert et al. (2021). Writing $\mathbf{q} = (\mathbf{R}, \mathbf{T})$, the equilibrium response is

$$d\mathbf{q} = -\mathcal{H}_q^{-1}\mathcal{H}_A \, d\mathbf{A}. \tag{25}$$

The distribution inherited at the first innovation and the previous bond return are fixed on impact. Idiosyncratic risk and borrowing constraints enter the stationary household problem nonlinearly. Individual policies retain the loading-specific response to current and expected future compensation, as well as the common equilibrium return and transfer responses. The resulting stationary response kernels can be obtained together with past aggregate shocks.

For the regression exercises, an aggregate innovation arrives in every annual period. Policies incorporate the realized past and the conditional AR(1) forecast; future realized innovations do not enter current choices. Persistent and transitory idiosyncratic shocks also continue each year. We use two histories of 201 annual dates, initialized at $A_{-1} = 0$, and observe outcomes every two years, giving 101 observations and 100 growth observations per household.

Pass-through for aggregate shocks. We estimate the consumption responses to aggregate income shocks using the same regression that we used in the empirical part. We first estimate the aggregate factor in the simulated log-income history of households. Let $\mu_{gt} = \mathbb{E}[\log y_{it} \mid \lambda_i = \lambda_g, \text{aggregate history}]$ and $\mu_t = \sum_g \omega_g \mu_{gt}$. We integrate over each type's household states to calculate these population conditional means. Using the known, centered loadings, define

$$F_t = \frac{\sum_g \omega_g \lambda_g (\mu_{gt} - \mu_t)}{\sum_g \omega_g \lambda_g^2}, \qquad \mu_{gt} - \mu_t = \lambda_g F_t + v_{gt}. \tag{26}$$

At biennial dates, we project F_t on its preceding observation to get $\rho_{F,2}$ as this auxiliary coefficient. Then the aggregate income innovation is

$$\widehat{\eta}_{it}^{A} = \lambda_i \left(F_t - \rho_{F,2} F_{t-2} \right). \tag{27}$$

The aggregate consumption projection is

$$\Delta_2 \log c_{it} = \mu_i^A + \delta_t^A + \phi_A \widehat{\eta}_{it}^{A} + u_{it}^A. \tag{28}$$

Because the regressor depends only on type and time period, its population coefficient can be computed from type-conditional mean log consumption growth with type and time effects and weights ω_g. Time effects remove the common consumption movement: the coefficient measures consumption differences associated with exposure.

Pass-through for idiosyncratic shocks. Subtract the same aggregate component and the factor:

$$s_{it} = \log y_{it} - \lambda_i F_t - \mu_t. \tag{29}$$

We fit the income process

$$s_{it} = \tilde{x}_{it} + \tilde{m}_{it}, \qquad \tilde{x}_{it} = \tilde{\rho}_{x,2} \tilde{x}_{i,t-2} + \tilde{\eta}_{it}^{x}, \tag{30}$$

with common persistent innovation and transitory variances across households. We obtain the idiosyncratic shocks using the same method as in the empirical part. The idiosyncratic response comes from the separate regression

$$\Delta_2 \log c_{it} = \mu_i^x + \delta_t^x + \phi_x \widehat{\tilde{\eta}_{it}^{x}} + u_{it}^x. \tag{31}$$

Table 6 reports the estimated pass-throughs from the model simulation and compares them to the data counterparts. The responses are 0.519 for aggregate shocks and 0.119 for idiosyncratic shocks. This full-information benchmark predicts an aggregate exposure re-

Table 6: Consumption Pass-throughs: Data and Model

Response	Data	Model
A. Gross labor earnings		
Aggregate, ϕ_A	0.126	0.519
	(0.088)	(0.050)
Idiosyncratic, ϕ_x	0.118	0.119
	(0.030)	(<0.001)

Data reproduce the baseline separate regressions in Table 2; parentheses report full-estimation household-bootstrap standard errors. The model column shows estimates from separate regressions from independent samples of 20,000 households per history, each with 100 biennial growth observations. The model columns use known fixed loadings and household and time effects. Standard errors below 0.0005 are displayed as < 0.001.

sponse about 4.4 times the idiosyncratic response. This is in stark contrast to what we found in the data. In the data, the ratio is close to 1.

3.5 Discussion

The measurement in the model and the empirical estimates share the same biennial consumption growth along with household- and time-effects. They differ in factor recovery: the model uses known loadings and exact conditional income means, whereas the empirical factor is estimated from observed household incomes. Model panels are long and balanced; empirical histories are short and unbalanced. The full-information benchmark matches the targeted idiosyncratic response but predicts a substantially larger aggregate exposure response.

The full-information benchmark predicts an aggregate exposure response approximately 4.4 times the idiosyncratic response. The empirical point estimates are much closer. This contrast motivates an investigation of information frictions as a possible explanation. In the next section, we develop intuition with a simple two-period model of a utility-maximizing household. The model admits a closed-form solution for the consumption function, thereby allowing a clear exposition of key mechanisms.

4 Why Full Information Fails? A Simple Two-Period Model

Before turning to the quantitative heterogeneous-agent model, we build intuition for the key mechanism using a simple two-period partial-equilibrium framework with CARA preferences. The analytical tractability of this setting allows us to derive the consumption function in closed form and to state precisely how incomplete information compresses the ratio ϕ_A/ϕ_x relative to the full-information benchmark. The mechanism isolated here, incomplete information, selectively attenuates the aggregate future-income channel while leaving the idiosyncratic channel unchanged, carries over intact to the quantitative model of Section 3.4.

4.1 Model Setup

Households. There is a unit continuum of households indexed by $i \in [0, 1]$. The economy lasts two periods, $t \in \{0, 1\}$. Household i has CARA preferences over consumption:

$$u(c) = -\frac{1}{\gamma}\, e^{-\gamma c}, \qquad \gamma > 0. \tag{32}$$

The household maximizes expected lifetime utility $u(c_{i0}) + \beta\, \mathbb{E}_0[u(c_{i1})]$ subject to the sequential budget constraints

$$c_{i0} + \frac{b_{i1}}{1+r} = b_{i0} + y_{i0}, \tag{33}$$

$$c_{i1} = (1+r)\, b_{i1} + y_{i1}, \tag{34}$$

where b_{it} is the bond position, r is the fixed partial-equilibrium interest rate, and y_{it} is labor income in period t. We set $\beta = 1$ and $r = 0$ (the no-discounting case) throughout for exposition.

Income Log labor income follows the same decomposition as in the quantitative model. In linearized form around the deterministic steady state:

$$y_{it} = (\lambda_i)\, A_t + x_{it}, \tag{35}$$

where A_t is the persistent aggregate component, x_{it} is the persistent idiosyncratic component, λ_i is a permanent household-specific loading on the aggregate. The two components evolve as

$$A_1 = \rho_A\, A_0 + \varepsilon_1^A, \qquad \varepsilon_1^A \sim \mathcal{N}(0, \sigma_A^2), \tag{36}$$

$$x_{i1} = \rho_x\, x_{i0} + \varepsilon_{i1}^x, \qquad \varepsilon_{i1}^x \sim \mathcal{N}(0, \sigma_x^2), \tag{37}$$

with $\rho_A > \rho_x > 0$, so the aggregate component is more persistent than the idiosyncratic component. The initial state A_0 has a prior distribution $A_0 \sim \mathcal{N}(0, \sigma_A^2)$, using the same innovation variance as in (36). The loadings λ_i take values in $\{\lambda_1, \ldots, \lambda_K\}$ with probability weights $\{\omega_k\}$.

Information structure. Under *full information* (FI), $\sigma_m = 0$ and both A_0 and x_{i0} are observed directly. Under *incomplete information* (II), A_0 is latent and $\sigma_m > 0$. Each household observes x_{i0} directly, but not A_0. Given x_{i0}, the household constructs the residual income signal

$$\tilde{y}_{i0} \equiv y_{i0} - x_{i0} = (\lambda_i)\, A_0 + m_{i0}, \qquad m_{i0} \sim \mathcal{N}(0, \sigma_m^2), \tag{38}$$

where m_{i0} is transitory noise, independent across households and of all other shocks. The residual $\tilde{y}_{i0}$ is the only signal about A_0. The household's information set at $t = 0$ is

$$\mathcal{I}_{i0} = \sigma(\tilde{y}_{i0},\, x_{i0})\,. \tag{39}$$

4.2 Consumption under Full Information

Combining the budget constraints (33)–(34) with $\beta = 1$ and $r = 0$ gives a two-period lifetime budget. With CARA utility and Gaussian income, the Euler equation pins down a closed-form consumption function.

Proposition 1 (Full-information consumption function)**.** *Under full information, the period-*0 *consumption of household* i *is*

$$c_{i0}^{\mathrm{FI}} = \frac{\lambda_i(1+\rho_A)}{2} A_0 + \frac{1+\rho_x}{2} x_{i0} - \frac{\gamma}{4}\left[\lambda_i^2\sigma_A^2 + \sigma_x^2\right]. \tag{40}$$

A unit innovation to A_0 raises current income by λ_i and the next period income via $\lambda_i\rho_A$; the present-value factor for the aggregate channel is therefore $(1+\rho_A)$. For the idiosyncratic component, the analogous factor is $(1+\rho_x)$. Since $\rho_A > \rho_x$, the consumption response to the aggregate shock exceeds that to the idiosyncratic shock (per unit of income loading), yielding $\phi_A > \phi_x$ under FI.

Bayesian Updating under Incomplete Information .

When A_0 is unobserved, the household uses signal (38) to form a posterior about A_0.

Given signal $\tilde{y}_{i0}$, the posterior distribution of A_0 is $\mathcal{N}(\hat{A}_{i0},\, P_i)$ with posterior variance and mean

$$P_i = \left[\frac{1}{\sigma_A^2} + \frac{\lambda_i^2}{\sigma_m^2}\right]^{-1}, \tag{41}$$

$$\hat{A}_{i0} = \frac{k_i}{\lambda_i}\,\tilde{y}_{i0}, \tag{42}$$

where the Kalman gain $k_i \in (0,1)$ is

$$k_i = \frac{\lambda_i^2\,\sigma_A^2}{\lambda_i^2\,\sigma_A^2 + \sigma_m^2}. \tag{43}$$

Substituting $\tilde{y}_{i0} = \lambda_i A_0 + m_{i0}$ into (42) gives $\hat{A}_{i0} = k_i A_0 + (k_i/\lambda_i)m_{i0}$, so $\partial\hat{A}_{i0}/\partial A_0 = k_i < 1$. The gain k_i is increasing in the signal-to-noise ratio λ_i^2/σ_m^2: households with a larger total income loading receive a more informative signal and thus track A_0 more closely. As $\sigma_m \to 0$, $k_i \to 1$ and FI is recovered; as $\sigma_m \to \infty$, $k_i \to 0$ and households learn nothing from the signal.

4.3 Consumption under Incomplete Information

Under incomplete information, the household replaces the unobserved A_0 with its posterior mean $\hat{A}_{i0}$ when forming expectations about future income.

Proposition 2 (Incomplete-information consumption function)**.** *Under incomplete information, the period-0 consumption of household i is*

$$c_{i0}^{\mathrm{II}} = \frac{\lambda_i(1+\rho_A k_i)}{2} A_0 + \frac{1+\rho_x}{2} x_{i0} + \frac{\rho_A k_i}{2} m_{i0} - \frac{\gamma}{4} \operatorname{Var}_0^{\mathrm{II}}[y_{i1}], \tag{44}$$

where $\operatorname{Var}_0^{\mathrm{II}}[y_{i1}] = \lambda_i^2(\rho_A^2 P_i + \sigma_A^2) + \sigma_x^2$, and m_{i0} is the private signal noise.

Comparing Propositions 1 and 2, incomplete information leaves the coefficient on x_{i0} unchanged at $(1+\rho_x)/2$, but replaces ρ_A with $\rho_A k_i < \rho_A$ in the coefficient on A_0. Because the household perceives only a fraction k_i of the true aggregate shock, it expects only $\lambda_i \rho_A k_i A_0$ in future income rather than $\lambda_i \rho_A A_0$. A fraction $1 - k_i$ of the innovation is attributed to noise.

Pass-Through Statistics and the Main Result We define the pass-through coefficients as:

$$\phi_A \equiv \frac{\operatorname{Cov}_i\big(\Delta c_{i0},\ \lambda_i \varepsilon^A\big)}{\operatorname{Var}_i(\lambda_i \varepsilon^A)}, \qquad \phi_x \equiv \frac{\operatorname{Cov}_i(\Delta c_{i0},\ \varepsilon_i^x)}{\operatorname{Var}_i(\varepsilon_i^x)}, \tag{45}$$

where $\varepsilon^A = A_0$ is the aggregate innovation, $\varepsilon_i^x = x_{i0}$ is the household-specific innovation, and cross-sectional variation in ϕ_A is driven by λ_i heterogeneity.

Proposition 3 (Incomplete information compresses ϕ_A/ϕ_x)**.** *Under full information,*

$$\phi_A^{\mathrm{FI}} = \frac{1+\rho_A}{2}, \qquad \phi_x = \frac{1+\rho_x}{2}, \qquad \frac{\phi_A^{\mathrm{FI}}}{\phi_x} = \frac{1+\rho_A}{1+\rho_x}. \tag{46}$$

Under incomplete information,

$$\phi_A^{\mathrm{II}} = \frac{1+\rho_A \bar{k}}{2}, \qquad \frac{\phi_A^{\mathrm{II}}}{\phi_x} = \frac{1+\rho_A \bar{k}}{1+\rho_x} < \frac{1+\rho_A}{1+\rho_x} = \frac{\phi_A^{\mathrm{FI}}}{\phi_x}, \tag{47}$$

where $\bar{k} = E[\lambda^2 k] \,/\, E[\lambda^2] \in (0,1)$ is the type-weighted average Kalman gain. The pass-through ϕ_x is the same under FI and II. Moreover:

(i) *As* $\sigma_m \to 0$ (full information), $\bar{k} \to 1$ *and* $\phi_A^{\text{II}}/\phi_x \to \phi_A^{\text{FI}}/\phi_x$.

(ii) *As* $\sigma_m \to \infty$ (pure noise), $\bar{k} \to 0$ *and* $\phi_A^{\text{II}}/\phi_x \to 1/(1+\rho_x)$.

(iii) *There exists a unique* $\sigma_m^* > 0$ *such that* $\phi_A^{\text{II}}/\phi_x = 1$, *achieved when* $\bar{k} = \rho_x/\rho_A$.

Why the aggregate channel is attenuated. Under incomplete information, the household attributes only a fraction k_i of the true shock to the aggregate component, so the expected future aggregate income is $\lambda_i \rho_A k_i A_0$ rather than $\lambda_i \rho_A A_0$. Therefore, the reduction in the present value of the income shock relative to the full information benchmark. This pushes the ratio ϕ_A/ϕ_x closer to 1. Moreover, the extent of the shock attributed to the aggregate component depends on the signal noise. With a noisier signal or a less informative signal, much of the movement in the residual income is attributed to noise rather than the aggregate component.

Connection to the quantitative model. The CARA two-period model isolates the core mechanism in its starkest form. We next turn to a quantitative GE model with two realistic features. Since we are interested in matching average pass-throughs, we embed household heterogeneity à la Bewley-Aiyagari-Huggett. We also allow financial frictions in the form of borrowing constraints. We solve this model in general equilibrium, accounting for GE effects of changes in the aggregate component of income on consumption responses. The results in the following section show that a full-information consumption-savings model still fails to match the cross-sectional empirical pattern we document, whereas a model with information frictions goes a long way toward reproducing the empirical facts. We also show the aggregate demand implications of introducing information frictions in a standard consumption-savings model.

5 Quantitative Incomplete-Information Economy

We now embed the signal-extraction problem from Section 4 into the same quantitative Bewley–Huggett economy. Relative to the full-information benchmark in Section 3, preferences, technology, market structure, and the steady-state computational backbone are un-

changed. What changes is the household's information set and, therefore, the perceived sequence of future income and interest-rate movements following an aggregate shock.

5.1 Information structure

Under *incomplete information* (II), $\sigma_m > 0$ and A_t is latent. Households observe x_{it} directly, and they observe total income y_{it}. Given knowledge of x_{it}, they compute the residual earnings signal

$$\tilde{y}_{it} \equiv \log y_{it} - x_{it} = \log w_t + \kappa_i A_t + m_{it} = \underbrace{(1+\kappa_i)}_{\text{Mapped to } \lambda_i} A_t + m_{it}, \tag{48}$$

where the second equality uses $\log w_t \approx A_t$. Households do not observe A_t but do know exactly the realization of x_{it}. By virtue of the noise component m_{it}, the households cannot back out A_t from their realized income. We assume that the residual $\tilde{y}_{it}$ is their only signal about A_t. This residual is a noisy scalar signal about A_t with type-specific loading $1 + \kappa_i \equiv \lambda_i$ and noise m_{it}. Households use this signal to update their beliefs about A_t each period.

Households do not receive any more information about the future path of equilibrium prices as signals to back out A_t; the filtering problem is based solely on income observations. However, once a household has formed its posterior $\hat{A}_{it}$ about the current aggregate state, we assume that it applies the correct general-equilibrium mapping from A to prices when projecting future interest rates. That is, households know the equilibrium pricing function $r(\cdot)$ and use it to translate their belief about A_0 into an expected path of r_t. The household's information set at time t is therefore

$$\mathcal{I}_{it} = \sigma\Big(\{\tilde{y}_{i\tau},\, x_{i\tau}\}_{\tau \le t}\Big). \tag{49}$$

Given the signal (48) and the AR(1) law of motion (14), households update beliefs via the Kalman filter. Let P_{it} denote household i's posterior variance and k_{it} the corresponding

Kalman gain. The posterior variance evolves as

$$\frac{1}{P_{it}} = \frac{1}{\rho_A^2 P_{i,t-1} + \sigma_A^2} + \frac{(\lambda_i)^2}{\sigma_m^2}, \tag{50}$$

and the posterior mean updates according to

$$\hat{A}_{it} = (1 - k_{it})\, \rho_A\, \hat{A}_{i,t-1} + \frac{k_{it}}{\lambda_i}\, \tilde{y}_{it}, \tag{51}$$

where $k_{it} = 1 - P_{it}/(\rho_A^2 P_{i,t-1} + \sigma_A^2)$. The posterior mean is thus a weighted average of the prior forecast and the new signal, with the weight on the signal increasing in the type-specific signal-to-noise ratio $(\lambda_i)^2/\sigma_m^2$.

5.2 Recursive household problem

Under incomplete information, A_t is not observed. Instead, the household carries its posterior mean $\hat{A}_{it}$ as a state. After observing x_{it} and the signal $\tilde{y}_{it}$, the household updates $\hat{A}_{it}$ via (51) and then makes its consumption-saving decision. The individual state after updating beliefs is $(b_{it},\, x_{it},\, \hat{A}_{it},\, \lambda_i)$. The technically complete aggregate state under II would be the joint distribution over $(\hat{A}_i,\, b_i,\, x_i,\, \lambda_i)$, since belief heterogeneity across households affects their savings decisions and hence aggregate bond demand. However, the noise m_{it} is drawn independently of $(x_{it},\, b_{it})$ across households, the cross-sectional distribution of posterior means $\hat{A}_{it}$ is independent of the joint distribution of $(b_i,\, x_i)$ at each date. This independence allows the aggregate state to be decomposed: the distribution Γ_t over $(b_i,\, x_i,\, \lambda_i)$ evolves as under FI, and the aggregate effect of beliefs on prices can be summarized by the type-specific average posterior $\bar{A}_{j,t} = \mathbb{E}[\hat{A}_{it} \,|\, \lambda_i = \lambda_j]$. In solving the model, we assume that households project the future path of interest rate using their own posterior $\hat{A}_{it}$; Note that owing to the independence mean zero of m_{it} over a large number of households of a given type j, the average belief maps cleanly as a linear function of realized A_t. Therefore, we have a clean mapping from realized A_t to r_t. Households apply the same GE mapping on their posterior to project r_t forward. The continuation values depend on the expected evolution of beliefs. The Bellman

equation is

$$V^{\text{II}}(b, x, \hat{A}, \lambda; \Gamma) = \max_{c,\, b' \geq \underline{b}} \left\{ \log c + \beta\, \mathbb{E}\Big[V^{\text{II}}(b', x', \hat{A}', \lambda; \Gamma') \;\Big|\; x,\ \hat{A} \Big] \right\} \tag{52}$$

subject to $c + b' = (1 + r(\Gamma))b + y$, where y is the realized income observed by the household.

Comparing (21) and (52), the two problems are structurally identical: a standard consumption-saving problem with one asset, borrowing limit, and an exogenous income process. The sole difference is in the state vector: A_t under FI is replaced by the posterior mean $\hat{A}_{it}$ under II. The independence of $\hat{A}_{it}$ from $(x_{it},\, b_{it})$ means that the distribution Γ_t and hence equilibrium prices can be computed using the same infrastructure in both environments. The informational friction does not require solving a fundamentally different model; it amounts to replacing the true aggregate state with the household's filtered estimate of it.

5.3 Solution under Incomplete Information

Incomplete-information solution. As noted above, the FI and II models share the same steady state, policy functions, and stationary distribution. We have a type-specific Kalman gain at k_j^* that we initialize for each type j. What changes relative to FI is the income path each type expects at $t = 0$. By the law of iterated expectations, household j's expectation at $t = 0$ of the aggregate income component s periods ahead is

$$\mathbb{E}_0\big[(\lambda_j) A_s \,\big|\, \mathcal{I}_{j,0}\big] = (\lambda_j)\, k_j^*\, \rho_A^s\, A_0, \tag{53}$$

attenuated relative to the FI path by the factor $k_j^* \in (0, 1)$. Substituting the perceived multiplier $(\lambda_j) k_j^* \rho_A^s$ in place of $(\lambda_j)\rho_A^s$ which delivers $\phi_A^{\text{II,PE}}$.

It is worth emphasizing what (53) is and is not. It is the household's current expectation of future income, the relevant object for today's consumption decision, formed using only information available at $t = 0$. It is not the path of future conditional expectations as the filter processes additional signals at $t = 1, 2, \ldots$; that path, which would incorporate expected future belief revisions, is not known to the household today. Future signals will arrive and shift beliefs, but at $t = 0$, their expected incremental contribution is zero by the martingale

property of the Kalman innovations. The result is the simple geometric decay in (53): the household discounts the persistence of the aggregate shock by the factor k_j^* that captures how much weight it places on the current signal relative to the prior.

How the II variant is solved with SSJ. The computational advantage of the SSJ approach is that the incomplete information variant can be solved by reusing the full-information Jacobians. In the FI model, the household block Jacobian maps sequences of aggregate inputs, income, and interest rates into consumption and savings responses. Under II, the household's problem is structurally the same; what changes is not the response to a unit input, but the sequence of inputs the household perceives. In particular, households react to their posterior mean $\hat{A}$ rather than the true aggregate state A.

This means that the FI Jacobian can be composed with a simple mapping from true aggregate shocks to perceived aggregate inputs. For type j, the perceived income sequence is attenuated by k_j^* at every horizon, so the FI income multiplier $(\lambda_j)\rho_A^s$ is replaced by $(\lambda_j)k_j^*\rho_A^s$. No additional backward iteration on household policy functions is required for this step: we use the same policy-function derivatives computed under FI and feed in the perceived income path as the relevant input sequence.

A second input sequence that changes under II is the path of equilibrium interest rates. Ideally, households need to forecast the path of dr_t. We assume that they are somehow given the correct path $dr^I I$ in this environment. While this can (and ideally should) be relaxed, we study it as the first case.

Thus, the household perceives a correct r path and an attenuated income path.

The remaining step is to determine the actual GE path dr^{II}. Because the noise m_{it} is i.i.d. across households and independent of (x_{it}, b_{it}), the cross-sectional distribution of beliefs evolves separately from the joint cross-sectional distribution of assets and idiosyncratic states. This allows aggregation by household type. That further simplifies the computation.

Since the backward EGM step is linear in the input sequences, the FI Jacobian can again be reused. Now, we are interested in the time t consumption response. For the income channel, the response of consumption (weighted by group) is calculated just as in the full information case, with a different input sequence $\mathrm{irf}_j(s) = k_j^*(\lambda_j)\rho_A^s$ for each type. Note that the

Kalman gain k_j^* attenuates the full magnitude of the shock to A_t. Now, for type-j r-channel MPC at time 0 is:

$$\delta_{II}^r c^j(x,a) = \sum_{s=0}^{T-1} \text{irf}_j(s) \cdot \frac{\partial c_0^j(x,a)}{\partial \, \mathrm{r}_s^{II}},$$

Aggregating over the stationary distribution at 0:

$$\psi_j^{r^{II}} \equiv \sum_{x,a} D^*\left(x,a \mid \lambda_j\right) \frac{\delta_{II}^r c^j(x,a)}{c^*(x,a)},$$

Given that the steady state consumption policy function is the same regardless of type j, $\psi_j^{r^{II}} = \psi^{r^{II}}$ for all j.

So the GE pass-through under II becomes

$$\phi_A^{\text{II,GE}} = \phi_A^{\text{II,PE}} + \psi^{r^{II}} \frac{\mathbb{E}[\lambda_j]}{\mathbb{E}[\lambda_j^2]} \tag{54}$$

The key point is that solving the II model does not require a new household block. The same SSJ objects from the FI benchmark are reused, with the true aggregate shock sequence replaced by the sequence perceived by households. This is the main computational reason the incomplete-information variant remains tractable.

The incomplete-information economy uses the same baseline calibration as in Section 3. The only additional object that matters for the information structure is the signal-noise parameter σ_m, which governs the quality of household information about the aggregate component. We set that to $\sigma_m = 0.45$, in line with our estimates.

5.4 Quantitative Results

Table 7 reports pass-through statistics under four specifications: FI and II, each in partial equilibrium (PE, interest rate fixed) and general equilibrium (GE, interest rate adjusts to clear the bond market).

Incomplete-information attenuation. Incomplete information attenuates the perceived size of the aggregate shock, thereby substantially compressing the discounted present value of

Table 7: Pass-Through Statistics: ϕ_A and ϕ_x

	FI-PE	FI-GE	II-PE	II-GE
	(1)	(2)	(3)	(4)
ϕ_A	0.78	0.79	0.24	0.20
ϕ_A/ϕ_x	4.61	4.67	1.40	1.12

Notes: PE holds r fixed at steady state. GE allows r to clear the bond market after a $1\sigma_A$ aggregate shock. I-GE assumes households observe the equilibrium interest rate but learn about aggregate productivity A_0 from their private income signals (via a Kalman filter). The equilibrium r-path under II is determined by bond-market clearing against the perceived income stream. Calibrated at $\beta = 0.95621$, $r_{ss} = 4.41\%$ annually, $\sigma_m = 0.34$, mass at borrowing constraint $= 2.8\%$.

the aggregate income path. As a result, $\phi_A^{\text{II,GE}}/\phi_x$ falls to 1.12, a substantial compression from around 4.6 under FI. Note that this occurs without targeting σ_m to hit $\phi_A/\phi_x \approx 1$. We use the estimated value of σ_m and obtain $\phi_A/\phi_x \approx 1.12$ as an untargeted moment.

A key finding is that the attenuation under II operates primarily through the income channel, not the r-channel. The actual GE r-path is nearly unchanged between FI and II. This is partly because we assume a known interest rate path even under incomplete information. and this path rapidly approaches the full-information path (the learning about A as time progresses occurs rapidly in this model).

Aggregate Demand Response Under Information Frictions. Now, we study the implications of information frictions for aggregate demand response to aggregate income shocks such as A_t. Aggregate consumption impulse response function (IRF), the response of total consumption $C_t = \int c_{it}\, di$ to a unit innovation $\varepsilon_0^A = 1$ captures the aggregate demand implications in our model.

Under FI-GE, the aggregate IRF follows directly from the Jacobians J in PE:

$$\frac{dC_t^{\text{FI,GE}}}{d\varepsilon_0^A} = \sum_{s=0}^{T-1} J^{C,Agg}(t,s)\,\mathbb{E}[\lambda]\,\rho_A^s + \sum_{s=0}^{T-1} J^{C,r}(t,s)\,dr_s^{\text{FI}}, \tag{55}$$

where $J^{C,Agg}$ is the household Jacobian mapping changes in Aggregate shock sequence to

change in consumption and $J^{C,r}$ is the household Jacobian mapping sequence of interest rate changes to consumption changes. Note that dr_s^{FI} is the realized GE interest-rate response at horizon s.

Under II, households gradually learn about A_0 from their private income signals. We assume that the equilibrium interest rate r_t is publicly observable, so there are *no forecast revisions on the interest rate*: the r-channel transmits the equilibrium rate fully to all households, without attenuation. Only the income channel is subject to learning.

This distinction is consequential for how beliefs enter the GE equilibrium. Let $\Delta\Theta_\tau = \Theta_\tau - \Theta_{\tau-1}$ (with $\Theta_{-1} \equiv 0$) denote the increment to the aggregate belief attenuation sequence, where $\theta_{j,t}$ is the fraction of A_0 that type j has learned by date t. At date τ, households update their collective beliefs by $\Delta\Theta_\tau$ and revise their savings plans for all future dates $t \geq \tau$. Define the *perceived aggregate path* as the convolution of these incremental belief updates with the true shock path:

$$A_t^{\text{II}} \;=\; \sum_{\tau=0}^{t} \Delta\Theta_\tau \; A_{t-\tau}^{\text{FI}}. \tag{56}$$

By construction, $A_0^{\text{II}} = \Theta_0 \, \sigma_A$, so on impact the perceived aggregate shock is attenuated by Θ_0.

The aggregate consumption IRF under II-GE then follows from the same GE Jacobian $J^{C,Agg}$ manipulated to account for the fact that forecast revisions at t only induce response after $s \geq t$ and not before (See Auclert et al. (2020), Guerreiro (2023) for details on this).

$$d\mathbf{C} = [\mathcal{J}^{C,Agg} \cdot A_0^{II} + \sum_{t\geq 1} \mathcal{J}_{t,LT}^{C,Agg} \cdot \underbrace{\left(A_t^{II}\right)}_{\text{Forecast Revision at time } t}] + \sum_{t=0}^{T-1} \mathcal{J}^{C,r} \cdot dr_t^{II},$$

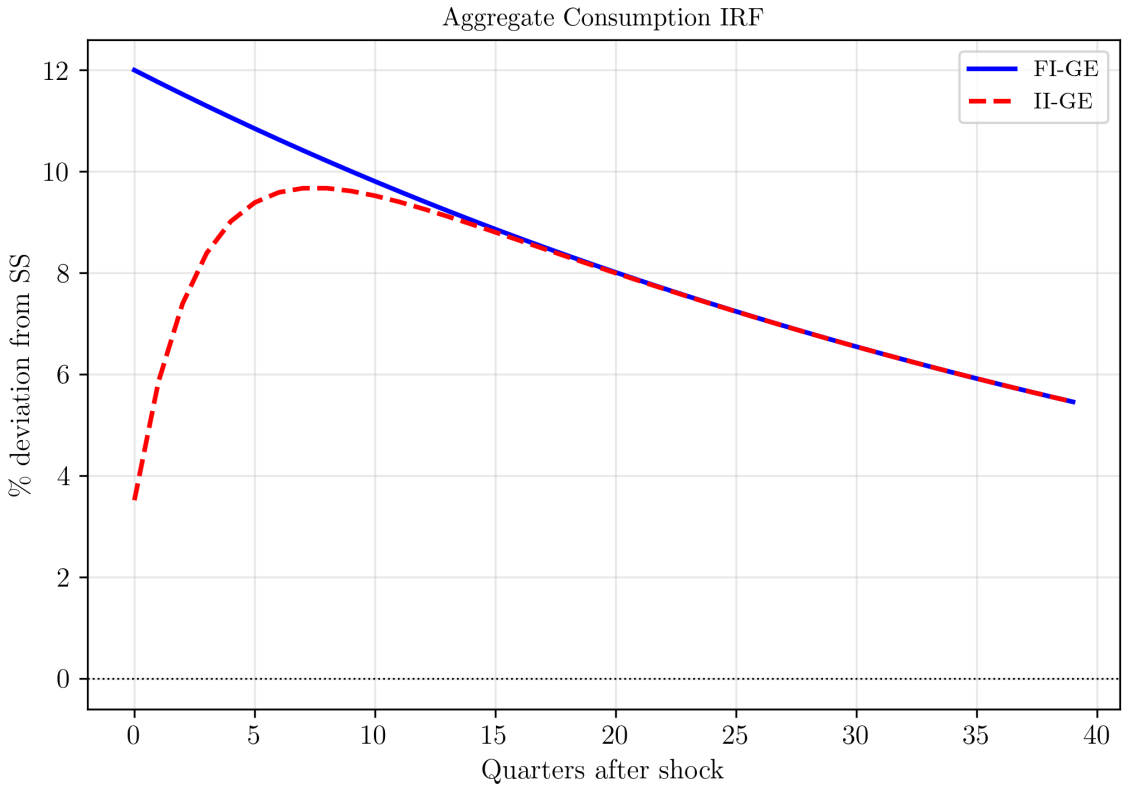
where

$$\mathcal{J}_{t,LT}^{C,Agg} \equiv \begin{bmatrix} 0 & \mathbf{0}_t' \\ \mathbf{0}_t' & \mathcal{J}^{C,Agg} \end{bmatrix}. \tag{57}$$

Note that, by assumption, we don't have any forecast revisions for r_t^{II}, since it is taken to be given.

The impact attenuation is $\Theta_0 = 0.293$. The FI-GE impact response is $dC_0^{\text{FI,GE}}/C = 12.0\%$,

Figure 4: Impulse Response of Aggregate Consumption



Notes: This figure shows the impulse response of aggregate consumption in both full information and incomplete information settings.

while the II-GE impact response is 3.5%—about 29% of the FI level. Households that cannot distinguish persistent aggregate innovations from transitory idiosyncratic noise consume far less on impact than under full information. As beliefs update, Θ_t rises rapidly, and the consumption under incomplete information rises to catch up with the full information level.

6 Conclusion

This paper documents a simple but striking fact: households respond almost equally to aggregate and idiosyncratic income shocks, even though the aggregate component of income is far more persistent. Under the standard full-information permanent-income logic, this should not happen. A more persistent shock should induce a substantially larger consumption response. We show that this is a puzzle for a standard incomplete-markets benchmark. In

a quantitative full-information Huggett economy, households respond much more strongly to aggregate than to idiosyncratic shocks. This prediction survives in richer models with realistic wealth heterogeneity and borrowing constraints. We argue that incomplete information about the source of income fluctuations provides a natural resolution. Households observe total income, but not its decomposition into aggregate and idiosyncratic components, so they underreact to aggregate shocks because they do not fully internalize their persistence. Embedding the same friction in a quantitative heterogeneous-agent model compresses the gap between aggregate and idiosyncratic pass-through and brings the model much closer to the data.

Finally, because households view aggregate shocks as less informative about future income than they truly are, incomplete information dampens the response of aggregate consumption to those shocks. More broadly, our findings suggest that understanding consumption dynamics requires more than measuring the persistence of income shocks. It also requires understanding what households know about those shocks, and how they process that information.

References

Adams, Jonathan J. and Eugenio Rojas, "Household Consumption and Dispersed Information," *Journal of Monetary Economics*, 2024, *147*, 103652.

Angeletos, George-Marios and Chen Lian, "Confidence and the Propagation of Demand Shocks," *Review of Economic Studies*, 2022, *89* (3), 1085–1119.

_ **and Jennifer La'O**, "Sentiments," *Econometrica*, 2013, *81* (2), 739–779.

Arellano, Manuel, Richard Blundell, and Stéphane Bonhomme, "Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework," *Econometrica*, 2017, *85* (3), 693–734.

Auclert, Adrien, "Monetary Policy and the Redistribution Channel," *American Economic Review*, 2019, *109* (6), 2333–2367.

_ , **Bence Bardóczy, Matthew Rognlie, and Ludwig Straub**, "Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models," *Econometrica*, 2021, *89* (5), 2375–2408.

_ , **Matthew Rognlie, and Ludwig Straub**, "Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated HANK Model," Working Paper 26647, National Bureau of Economic Research January 2020.

Blundell, Richard, Luigi Pistaferri, and Ian Preston, "Consumption Inequality and Partial Insurance," *American Economic Review*, 2008, *98* (5), 1887–1921.

Braxton, J. Carter, Kyle Herkenhoff, Jonathan Rothbaum, and Lawrence Schmidt, "Changing Income Risk across the US Skill Distribution: Evidence from a Generalized Kalman Filter," *American Economic Review*, December 2025, *115* (12), 4438–75.

Caplin, Andrew, Victoria Gregory, Eungik Lee, Søren Leth-Petersen, and Johan Sæverud, "Subjective Earnings Risk," 2023. NBER Working Paper 31019.

Carroll, Christopher D., "The method of endogenous gridpoints for solving dynamic stochastic optimization problems," *Economics Letters*, 2006, *91* (3), 312–320.

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

Ghosh, Anisha and Alexandros Theloudis, "Consumption Partial Insurance in the Presence of Tail Income Risk," Papers 2306.13208, arXiv.org Jun 2023.

Guerreiro, Joao, "Belief disagreement and business cycles," *Manuscript, February*, 2023.

Guerrieri, Veronica and Guido Lorenzoni, "Credit Crises, Precautionary Savings, and the Liquidity Trap," *Quarterly Journal of Economics*, 2017, *132* (3), 1427–1467.

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

_ , "An empirical investigation of labor income processes," *Review of Economic Dynamics*, 2009, *12* (1), 58–79.

_ **and Anthony A. Smith**, "Inferring Labor Income Risk and Partial Insurance from Economic Choices," *Econometrica*, 2014, *82* (6), 2085–2129.

_ **, Serdar Ozkan, and Jae Song**, "The Nature of Countercyclical Income Risk," *Journal of Political Economy*, 2014, *122* (3), 621–660.

Hall, Robert E., "Stochastic Implications of the Life Cycle–Permanent Income Hypothesis: Theory and Evidence," *Journal of Political Economy*, 1978, *86* (6), 971–987.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante, "Consumption and Labor Supply with Partial Insurance: An Analytical Framework," *American Economic Review*, 2014, *104* (7), 2075–2126.

Jordà, Oscar and Alan M. Taylor, "Local Projections," *Journal of Economic Literature*, March 2025, *63* (1), 59–110.

Kaplan, Greg and Giovanni L. Violante, "How Much Consumption Insurance beyond Self-Insurance?," *American Economic Journal: Macroeconomics*, 2010, *2* (4), 53–87.

_ **and** _ , "A Model of the Consumption Response to Fiscal Stimulus Payments," *Econometrica*, 2014, *82* (4), 1199–1239.

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

Lorenzoni, Guido, "A Theory of Demand Shocks," *American Economic Review*, 2009, *99* (5), 2050–2084.

McKay, Alisdair and Ricardo Reis, "The Role of Automatic Stabilizers in the U.S. Business Cycle," *Econometrica*, 2016, *84* (1), 141–194.

Meghir, Costas and Luigi Pistaferri, "Income Variance Dynamics and Heterogeneity," *Econometrica*, 2004, *72* (1), 1–32.

_ **and** _ , "Earnings, Consumption and Life Cycle Choices," in Orley Ashenfelter and David Card, eds., *Handbook of Labor Economics*, Vol. 4B, Elsevier, 2011, pp. 773–854.

Pischke, Jörn-Steffen, "Individual Income, Incomplete Information, and Aggregate Consumption," *Econometrica*, 1995, *63* (4), 805–840.

Òscar Jordà, "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, March 2005, *95* (1), 161–182.

A Data and Estimation

In this appendix, we provide additional details on the data and estimation procedures underlying our empirical analysis. We first describe the data sources and sample construction, followed by the methodology for estimating the income process and recovering aggregate and idiosyncratic income shocks. We then discuss the estimation of consumption pass-through coefficients.

A.1 Sample Construction

We use household income and consumption data from the Panel Study of Income Dynamics (PSID) to estimate the income process and the consumption pass-through of different types of income shocks. The PSID is a longitudinal survey of U.S. households that began in 1968 with approximately 5,000 households. The survey was conducted annually until 1997 and transitioned to a biennial frequency beginning in 1999. Prior to 1999, consumption data were largely limited to food expenditures. Beginning in 1999, the PSID expanded its coverage to include a broader range of consumption categories. We therefore restrict our analysis to eleven biennial waves spanning 1999 to 2019, allowing us to maintain a consistent survey frequency and use more comprehensive measures of household consumption.

Income measure: We measure household labor income as the sum of pre-tax labor earnings of all household members, expressed in real terms. As a robustness check, we alternatively use the labor earnings of the household head. Following standard practice in PSID-based studies, we exclude observations from the Survey of Economic Opportunity (SEO) and Latino samples to address potential selection concerns arising from their sampling designs.

Consumption measure: Prior to 1999, the PSID collected expenditure information on a limited set of consumption categories, primarily food, housing, and childcare. Beginning in 1999, the survey substantially expanded its coverage to include categories accounting for approximately 70% of total expenditures measured in the Consumer Expenditure Survey (CEX). These categories include housing-related expenditures, utilities, transportation, education, and healthcare. The coverage was further expanded in 2005 to include additional categories such as clothing, household furnishings, and home repairs and maintenance. Since 2005, the

Table A1: Summary Statistics of PSID Sample

	Mean	SD	P25	P50	P75	N
Labor Income (log)	11.69	0.626	11.319	11.681	12.043	14819
Consumption (log)	10.09	0.485	9.768	10.055	10.373	14819
Age	47.37	9.238	40	47	55	14819
Education (# of years)	14.42	2.168	12	15	16	14819

Notes: This table shows summary statistics for the sample obtained from the PSID for the time period 1999-2019. The variable 'Labor Income' is the total household income earned through wages, the variable 'Consumption' is the total household consumption, and the variables 'Age' and 'Education' are the age and education of the household head, respectively.

PSID has covered nearly all major expenditure categories included in the CEX. For our baseline analysis, we restrict the consumption measure to categories consistently available since 1999 to ensure comparability across survey waves.[8]

Sample Selection: Our baseline sample consists of eleven biennial waves of the PSID from 1999 to 2019. We restrict the sample to households whose heads are between 30 and 65 years of age and exclude observations with annual household labor income below $5,000 or top-coded labor income. All income measures are expressed in 2019 U.S. dollars. We use information on years of schooling for the household head and spouse or partner, recorded on a scale from 0 to 17, along with demographic controls including race, age, and age polynomials for the household head and spouse or partner, family size, and state of residence. We drop observations with missing values for any of these variables. To ensure sufficient longitudinal information for estimating the income process, we further require each household to have at least five observed income waves. These restrictions yield a final unbalanced panel of 1,908 households and 14,819 observations.

A.2 Summary Statistics

Table A1 provides the summary statistics. The sample shows a mean log labor income of 11.69, a substantial standard deviation of 0.63, and a median log labor income of 11.68.

[8]We interpret our consumption measure as a proxy for nondurable consumption expenditures and services.

Table A2: Moments of Income and Consumption Residuals

Moment	Income	Consumption	Joint
Variance	0.295	0.123	
	(0.011)	(0.003)	
Skewness	0.085	-0.057	
	(0.125)	(0.035)	
Kurtosis	6.124	4.535	
	(0.439)	(0.113)	
Covariance			0.004
			(0.001)

Notes: This table shows the moments of residual income and residual consumption. Block bootstrap standard errors are reported in the parentheses.

The average of log consumption expenditure is 10.09, with a standard deviation of 0.48 and a median of 10.05. Since we selected the final sample of household heads aged 30-65, the sample mean is 47, and the standard deviation is about 9. The average years of schooling for the household head in the sample is around 14 years, and the median schooling is 15 years.

For our analysis, we residualize the log of labor income and the log of consumption on a rich set of observables, including third-order polynomials in age, dummies for race, region, race-year, region-year, family size, year of birth, number of children, and outside dependents. We interpret these residuals as unexplained income and consumption. Table A2 presents the empirical moments calculated using the residualized income and consumption. The variance of income is around 0.29, and the variance of consumption is around 0.12. The covariance between income and consumption residuals is 0.004 and significant, which determines the pass-through coefficients of income shocks to consumption. Also, note that the kurtosis is high in both the income and consumption residuals, indicating tail risk.

A.3 Estimation Details

As described in the main text, we model residualized log household income, z_{it}, as the sum of an aggregate component with heterogeneous household exposure, a persistent

idiosyncratic component, and a transitory component. Specifically, the income process is given by

$$\begin{aligned} z_{it} &= \lambda_i A_t + x_{it} + m_{it}, \\ A_t &= \rho_A A_{t-1} + \epsilon_t^A, \\ x_{it} &= \rho_x x_{i,t-1} + \epsilon_{it}^x, \end{aligned} \tag{A1}$$

where the innovations are normally distributed:

$$m_{it} \sim \mathcal{N}(0, \sigma_m^2), \qquad \epsilon_t^A \sim \mathcal{N}(0, \sigma_A^2), \qquad \epsilon_{it}^x \sim \mathcal{N}(0, \sigma_x^2). \tag{A2}$$

Households differ in their exposure to the common aggregate factor through the loading λ_i, while the persistence and variance parameters governing the idiosyncratic and transitory components, ρ_x, σ_x^2, and σ_m^2, are assumed to be common across households. We assume that all innovations are serially independent, that idiosyncratic and transitory innovations are independent across households, and that the three types of innovations are mutually independent. The initial idiosyncratic state is drawn from its stationary distribution. The transitory component may also capture measurement error, which is not separately identified from transitory income shocks.

An important identification consideration is that the aggregate factor A_t and its household-specific loadings λ_i are not separately identified without a normalization, as the factor model identifies only their product $\lambda_i A_t$. Consequently, while the aggregate persistence ρ_A is invariant to the normalization, the variance of aggregate innovations σ_A^2 depends on the scale of the factor. We therefore focus on the persistence of the aggregate component and the household-level aggregate innovations $\lambda_i \epsilon_t^A$, which are invariant to the factor normalization.

Since the PSID observations are biennial, the time index t denotes survey waves, and the persistence parameters ρ_A and ρ_x govern income dynamics at a two-year frequency. To facilitate comparisons with the literature, we report their annual equivalents, $\rho_A^{\text{ann}} = \sqrt{\rho_A}$ and $\rho_x^{\text{ann}} = \sqrt{\rho_x}$, assuming an underlying annual AR(1) process with positive persistence. The innovation variances σ_A^2 and σ_x^2, as well as the transitory variance σ_m^2, are maintained at their original wave-level frequency and are not annualized.

We first estimate the common aggregate factor and household-specific loadings using

uncentered rank-one principal component analysis (PCA). Specifically, we solve

$$(\widehat{A}, \widetilde{\lambda}) = \arg\min_{a,\ell} \sum_{i=1}^{N} \sum_{t=1}^{T} (\widetilde{z}_{it} - \ell_i a_t)^2, \qquad \text{subject to} \quad \frac{1}{N} \sum_{i=1}^{N} \ell_i^2 = 1, \tag{A3}$$

where $\widetilde{z}_{it}$ denotes residualized household income, with missing observations replaced by zero for the PCA estimation. The normalization fixes the scale of the factor and its loadings, while we resolve the sign indeterminacy by requiring the average household loading to be positive. This procedure yields an estimate of the aggregate factor, $\widehat{A}_t$, and preliminary household-specific loadings, $\widetilde{\lambda}_i$. The replacement of missing observations with zeros is restricted to the PCA step; subsequent income-process estimation and consumption regressions use only observed household records.

We estimate aggregate persistence separately from the idiosyncratic income process. Define the fitted first-stage common component as $\widetilde{g}_{it} = \widetilde{\lambda}_i \widehat{A}_t$. Its pooled persistence estimate is

$$\widehat{\rho}_A = \arg\min_{-1+\kappa \le r \le 1-\kappa} \sum_{i} \sum_{t=2}^{T} (\widetilde{g}_{it} - r\widetilde{g}_{i,t-1})^2, \qquad \kappa = 10^{-6}.$$

Because the time-invariant loadings enter both terms multiplicatively, the unrestricted solution is

$$\widehat{\rho}_A^{\mathrm{u}} = \frac{\sum_{t=2}^{T} \widehat{A}_{t-1} \widehat{A}_t}{\sum_{t=2}^{T} \widehat{A}_{t-1}^2}$$

Thus, aggregate persistence is identified directly from the time-series dynamics of the estimated common factor and does not depend on the scale or distribution of household-specific loadings. The constrained estimate equals the unrestricted solution whenever it lies within the stationary interval $[-1+\kappa, 1-\kappa]$ and otherwise equals the nearest boundary. Importantly, replacing the preliminary PCA loadings with their second-stage GLS estimates does not alter the estimated persistence parameter, provided the estimated factor is held fixed. Finally, we recover aggregate income innovations as

$$\widehat{\epsilon}_t^A = \widehat{A}_t - \widehat{\rho}_A \widehat{A}_{t-1}, \qquad t = 2, \ldots, T. \tag{A4}$$

These innovations are subsequently scaled by household-specific loadings to construct the

aggregate income shocks used in the consumption pass-through estimation. The estimated biennial persistence of the aggregate component is $\widehat{\rho}_A = 0.992$, corresponding to an annual persistence of $\widehat{\rho}_A^{\text{ann}} = 0.996$. The stationarity restriction does not bind at the estimated value. Although the estimation criterion pools information across households, aggregate persistence is identified from the time-series variation in the common factor, which contains only ten distinct calendar transitions. Thus, pooling across households does not provide additional independent observations of aggregate income dynamics. The estimated coefficient characterizes the persistence of the fitted common component, conditional on the recovered aggregate factor.

Having estimated the aggregate component, we remove its contribution to household income to isolate the remaining idiosyncratic variation: $\widehat{s}_{it} = z_{it} - \lambda_i A_t$. We model this residual variation as the sum of a persistent idiosyncratic component and a transitory component, using the following state-space representation:

$$\begin{aligned} \text{Measurement:} \quad & \widehat{s}_{it} = x_{it} + m_{it}, \qquad m_{it} \sim \mathcal{N}(0, \sigma_m^2), \\ \text{State:} \quad & x_{it} = \rho_x x_{i,t-1} + \epsilon_{it}^x \qquad \epsilon_{it}^x \sim \mathcal{N}(0, \sigma_x^2). \end{aligned} \tag{A5}$$

We estimate the parameters $\rho_x, \sigma_x^2, \sigma_m^2$, which are common across households, and recover the latent idiosyncratic states x_{it} using a Kalman filter combined with the Expectation-Maximization (EM) algorithm, following Braxton et al. (2025) closely. The algorithm iterates between two steps: the E-step computes the expected sufficient statistics of the latent states using Kalman smoothing, while the M-step updates the parameters by maximizing the expected complete-data log-likelihood. Iterating between these steps yields estimates of the persistence and shock variances of the idiosyncratic income process, along with the smoothed latent states used to recover idiosyncratic income innovations.

We initialize the latent state using its stationary distribution, $x_{i1} \sim \mathcal{N}(0, \sigma_x^2/(1-\rho_x^2))$, under the assumption $|\rho_x| < 1$. The Kalman filtering and smoothing procedures operate on the full biennial time grid, $t = 1, \ldots, T$, accommodating the unbalanced nature of the PSID panel. At dates for which household income is missing, the measurement update is omitted, while the latent state continues to evolve according to its transition equation. Thus, unlike

the initial PCA estimation, missing income observations are not replaced with zeros in the state-space estimation.

E-step: At iteration k, use the current parameter vector $\theta^{(k)}$. For readability, suppress the iteration superscript on the filter moments. Let $\mu_{it|r}$ and $P_{it|r}$ denote the conditional mean and variance of x_{it} given household i's observed adjusted income through date r. Initialize $\mu_{i1|0} = 0$ and $P_{i1|0} = \sigma_x^{2,(k)}/[1 - (\rho_x^{(k)})^2]$. For $t \geq 2$, the prediction step is

$$\mu_{it|t-1} = \rho_x^{(k)} \mu_{i,t-1|t-1},$$
$$P_{it|t-1} = (\rho_x^{(k)})^2 P_{i,t-1|t-1} + \sigma_x^{2,(k)}.$$

At an observed income date, define the measurement innovation and its variance as

$$\omega_{it} = \widehat{s}_{it} - \mu_{it|t-1}, \qquad F_{it} = P_{it|t-1} + \sigma_m^{2,(k)}.$$

The Kalman gain and measurement updates are

$$G_{it} = P_{it|t-1}/F_{it},$$
$$\mu_{it|t} = \mu_{it|t-1} + G_{it}\omega_{it}, \qquad P_{it|t} = (1 - G_{it})P_{it|t-1}.$$

When income is missing, skip the measurement update and retain $\mu_{it|t} = \mu_{it|t-1}$ and $P_{it|t} = P_{it|t-1}$. Starting from the filtered moments at T, run the backward smoother for $t = T - 1, \ldots, 1$:

$$J_{it} = \frac{\rho_x^{(k)} P_{it|t}}{P_{i,t+1|t}},$$
$$\mu_{it|T} = \mu_{it|t} + J_{it}(\mu_{i,t+1|T} - \mu_{i,t+1|t}),$$
$$P_{it|T} = P_{it|t} + J_{it}^2(P_{i,t+1|T} - P_{i,t+1|t}).$$

The adjacent-state conditional covariance is $P_{it,t-1|T} = J_{i,t-1}P_{it|T}$ for $t \geq 2$. The moments needed in the M-step are therefore

$$\begin{aligned}
\overline{x}_{it}^{(k)} &= \mu_{it|T}, \\
H_{it}^{(k)} &= \mathbb{E}_k[x_{it}^2 \mid \widehat{\mathbf{s}}_i] = P_{it|T} + \mu_{it|T}^2, \\
L_{it}^{(k)} &= \mathbb{E}_k[x_{it}x_{i,t-1} \mid \widehat{\mathbf{s}}_i] = P_{it,t-1|T} + \mu_{it|T}\mu_{i,t-1|T}.
\end{aligned}$$

Here conditioning on $\widehat{\mathbf{s}}_i$ means conditioning on its observed entries; $\mathbb{E}_k$ evaluates expectations at $\theta^{(k)}$.

M-step: The transitory variance update is

$$\sigma_m^{2,(k+1)} = \frac{1}{n_{\text{obs}}} \sum_{(i,t)} \left[(\widehat{s}_{it} - \overline{x}_{it}^{(k)})^2 + P_{it|T} \right].$$

The smoothed state variance $P_{it|T}$ is essential: squared deviations from the smoothed mean alone omit uncertainty about the latent state. Only observed income cells enter this measurement-variance update. For the state parameters, define

$$\begin{aligned}
B_k &= \sum_i H_{i1}^{(k)}, \\
D_k(r) &= \sum_i \sum_{t=2}^{T} \left[H_{it}^{(k)} - 2rL_{it}^{(k)} + r^2 H_{i,t-1}^{(k)} \right].
\end{aligned}$$

All latent transitions on the biennial grid enter $D_k(r)$, including transitions across dates with missing measurements. Stationary initialization contributes an additional parameter-dependent term. The expected complete-data log likelihood for the state parameters is

$$\mathcal{Q}_x(r, v \mid \theta^{(k)}) = -\frac{1}{2} \left\{ NT \log v - N \log(1 - r^2) + \frac{D_k(r) + (1 - r^2)B_k}{v} \right\}, \quad v = \sigma_x^2.$$

For a candidate persistence r, its maximizing innovation variance is

$$\mathcal{V}_x^{(k)}(r) = \frac{D_k(r) + (1 - r^2)B_k}{NT}.$$

Thus the exact stationary M-step reduces to a one-dimensional minimization, followed by a variance update:

$$\rho_x^{(k+1)} = \arg\min_{r\in\mathcal{R}} \left\{ NT \log \mathcal{V}_x^{(k)}(r) - N\log(1-r^2) \right\},$$
$$\sigma_x^{2,(k+1)} = \mathcal{V}_x^{(k)}(\rho_x^{(k+1)}).$$

Here $\mathcal{R}$ is a closed interval inside $(-1,1)$; the baseline persistence bounds would give $\mathcal{R} = [-0.995, 0.995]$. The transition update $\rho_x^{(k+1)} = \sum_{i,t\geq 2} L_{it}^{(k)} / \sum_{i,t\geq 2} H_{i,t-1}^{(k)}$ applies when the initial-state density is fixed independently of these parameters, and the bound is inactive.

Convergence: At each iteration, evaluate the observed-data working log likelihood using the filter innovations:

$$\ell_s(\theta) = -\frac{1}{2}\sum_{(i,t)} \left[\log(2\pi) + \log F_{it} + \frac{\omega_{it}^2}{F_{it}} \right].$$

Iterate until the likelihood and parameter changes stabilize, keeping variances positive and persistence inside $\mathcal{R}$. Multiple admissible starting values would be used to assess convergence to local solutions. After convergence, rerun the smoother at the final estimates and construct the estimated persistent idiosyncratic shocks

$$\widehat{\epsilon}_{it}^{x,\mathrm{EM}} = \widehat{x}_{it|T}^{\mathrm{EM}} - \widehat{\rho}_x^{\mathrm{EM}} \widehat{x}_{i,t-1|T}^{\mathrm{EM}}.$$

Estimation without the aggregate factor: To assess the role of the common component, we re-estimate the income process with $\lambda_i = 0$ for every household. Denoting parameters and states in this specification by a subscript or superscript zero, the model becomes

$$z_{it} = x_{it}^0 + m_{it}^0, \qquad m_{it}^0 \sim \mathcal{N}(0, \sigma_{m,0}^2),$$
$$x_{it}^0 = \rho_{x,0} x_{i,t-1}^0 + \epsilon_{it}^{x,0}, \qquad \epsilon_{it}^{x,0} \sim \mathcal{N}(0, \sigma_{x,0}^2).$$

We retain the original income sample and stationary initialization. With no factor loading to eliminate, we maximize the Gaussian likelihood of observed income directly, accounting for actual wave gaps. We re-estimate persistence and both variances, smooth the latent state, and

construct $\widehat{\epsilon}^{x,0}_{it} = \widehat{x}^0_{it} - \widehat{\rho}_{x,0}\widehat{x}^0_{i,t-1}$. Consumption growth is regressed on this innovation with the same household and wave effects, response cutoffs, and observations as the factor-included specification.

B Additional Empirical Results

B.1 Robustness

We conduct several robustness checks to assess whether the similarity in consumption pass-through to aggregate and idiosyncratic income shocks is sensitive to sample selection. Panel A of Table B1 examines the role of liquidity constraints. Since liquidity-constrained households may respond strongly to income shocks regardless of their persistence, we progressively exclude households with low income (bottom income quartile) and low wealth (bottom wealth decile), who are more likely to face borrowing constraints. Excluding low-income households yields pass-through coefficients of 0.104 and 0.122 for aggregate and idiosyncratic shocks, respectively. Further excluding low-wealth households produces estimates of 0.088 and 0.116. Although the aggregate pass-through declines slightly, the estimates remain relatively close, suggesting that the similarity in consumption responses is not driven solely by households more likely to face liquidity constraints.

Panel B examines the sensitivity of our results to the length of household income histories and the presence of outliers. Restricting the sample to households observed for at least seven or nine waves yields aggregate pass-through estimates of 0.105 and 0.161, respectively, compared with idiosyncratic pass-through estimates of 0.122 and 0.136. Similarly, trimming observations at the 1% level produces estimates of 0.142 and 0.118. Across these specifications, the estimated consumption responses remain broadly similar despite the substantial difference in the persistence of the underlying income shocks. Moreover, the estimated idiosyncratic pass-through is consistently higher when the aggregate factor is excluded from the income process, reinforcing the importance of accounting for aggregate income fluctuations. Overall, these robustness checks suggest that our main empirical finding is not driven by liquidity constraints, the length of observed income histories, or extreme observations.

Table B1: Robustness of Pass-Through Coefficients

	Aggregate	Idiosyncratic	No factor	N	HH
A. Robustness to liquidity constraints					
Baseline	0.126	0.118	0.188	12,543	1,907
	(0.088)	(0.030)	(0.036)		
Exclude low income households	0.104	0.122	0.173	9,484	1,430
	(0.099)	(0.032)	(0.042)		
Exclude low income and low wealth households	0.088	0.116	0.167	8,832	1,317
	(0.103)	(0.034)	(0.046)		
B. Robustness to Income histories and Outliers					
At least 7 waves	0.105	0.122	0.191	9,689	1,248
	(0.089)	(0.037)	(0.041)		
At least 9 waves	0.161	0.136	0.199	6,419	726
	(0.124)	(0.056)	(0.047)		
Trimming at 1%	0.142	0.118	0.185	12,175	1,901
	(0.095)	(0.032)	(0.040)		

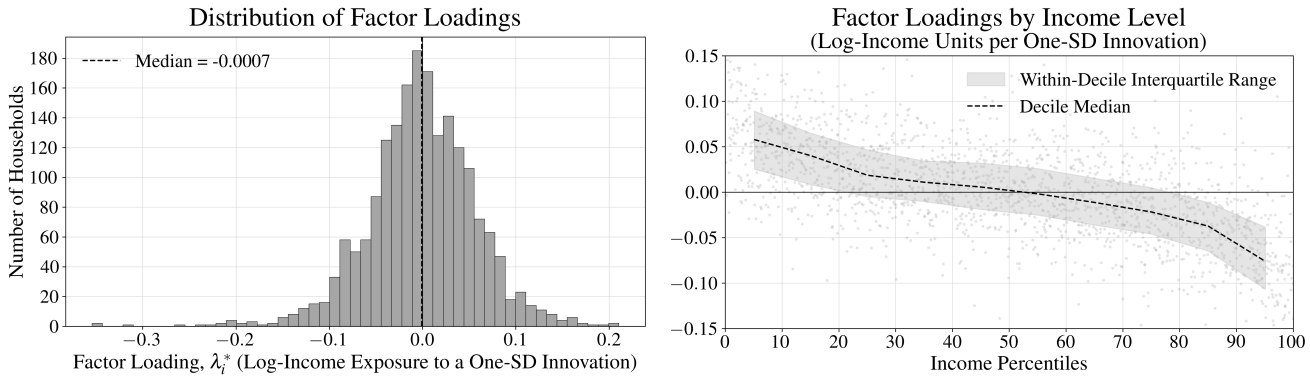
Notes: This table presents the robustness results for aggregate and idiosyncratic pass-throughs. Panel A shows the robustness of pass-through coefficients to excluding households with liquidity constraints, and Panel B shows the robustness to income histories and outliers. Low income and low wealth are used as proxies for liquidity constraints. Income shocks are winsorized at a 0.1% except in the 1% row. Block bootstrap standard errors are reported in the parentheses.

B.2 Factor Loadings and Heterogeneous Exposure

Figure B1 illustrates the heterogeneity in household exposure to aggregate income shocks. We normalize the estimated factor loadings by the standard deviation of aggregate innovations so that $\lambda_i^* = \lambda_i \sigma_A$ measures the change in household log income associated with a one-standard-deviation innovation to the aggregate factor. Panel (a) presents the distribution of estimated factor loadings across households. Although the distribution is concentrated around zero, with a median of approximately -0.0007, there is substantial heterogeneity in both the magnitude and direction of household exposure. While some households experience positive income innovations following a positive aggregate factor shock, others experience negative innovations, indicating that aggregate fluctuations need not affect all households uniformly.

Panel (b) examines how factor loadings vary across the household income distribution. The median loading declines systematically with household income, with positive loadings

Figure B1: Estimated Factor Loadings



(a) Distribution

(b) By Household Income

Notes: The figure presents the estimated household-specific factor loadings, normalized to measure changes in log household income associated with a one-standard-deviation innovation to the aggregate factor. Panel (a) shows the distribution of normalized loadings across households, with the dashed vertical line indicating the median. Panel (b) plots the estimated loadings against household income percentiles. The dashed line represents the median loading within each income decile, while the shaded region indicates the corresponding interquartile range.

concentrated among lower-income households and negative loadings becoming more prevalent among higher-income households. This pattern suggests that households at different points in the income distribution exhibit systematically different exposures to the common aggregate factor. Moreover, the substantial dispersion of loadings within income deciles indicates that aggregate income exposure also varies considerably among households with similar income levels. These findings underscore the importance of allowing heterogeneous household exposure to aggregate income shocks, as captured by λ_i, when estimating the income process and the subsequent consumption pass-through coefficients.

C Details on Full Information Benchmark Model

C.1 Solution Details

C.2 Model Impulse Response to Aggregate Shocks

For impulse responses, introduce one unanticipated aggregate innovation of size 0.002 at date zero and set subsequent aggregate innovations to zero. Aggregate productivity follows $A_h = 0.002\rho_A^h$. Idiosyncratic shocks continue, with their effects integrated over house-

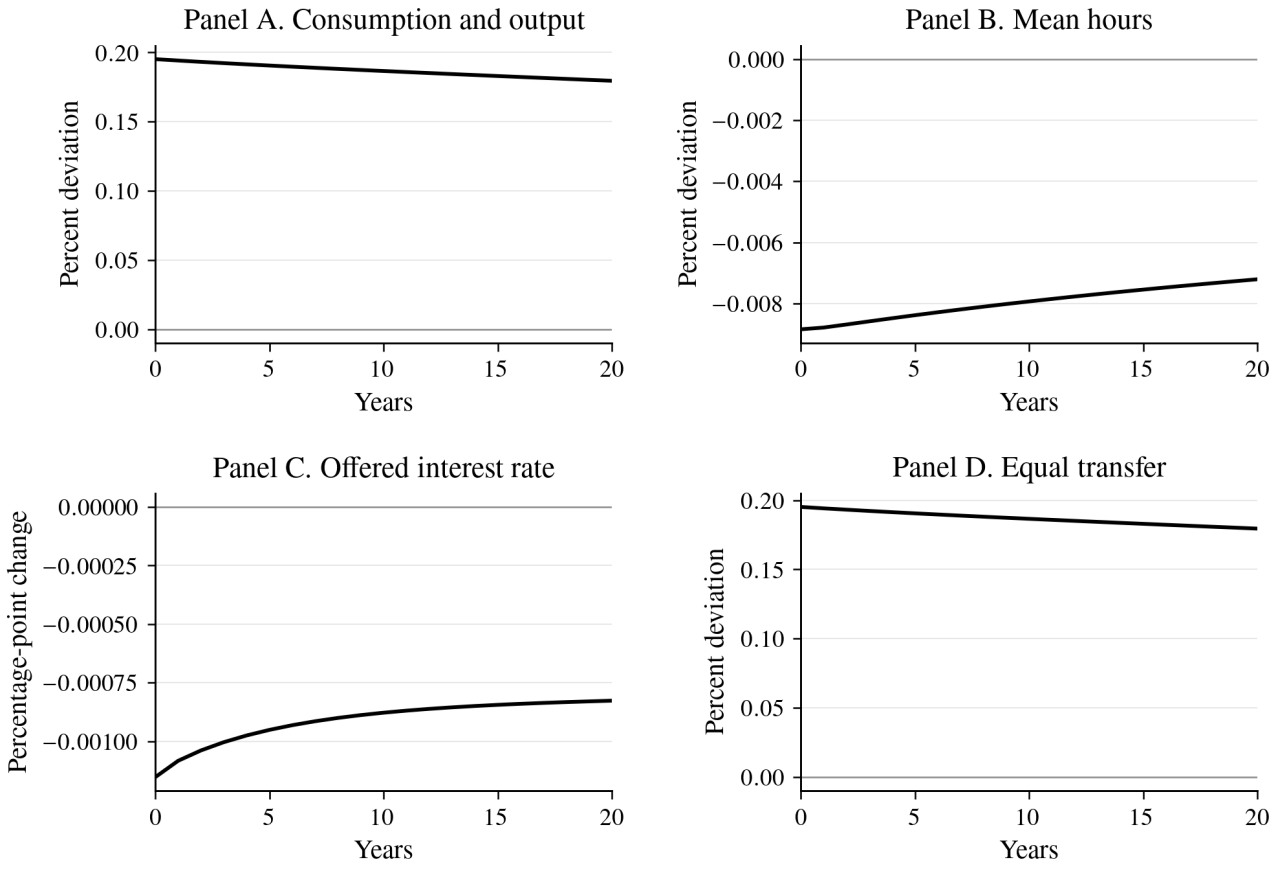


Figure C1: Aggregate responses to an aggregate productivity innovation

Notes: Responses to a positive innovation of 0.002 at date zero; subsequent aggregate innovations are zero. Panels A, B and D report percent deviations from the aggregate-shock-free path. Panel C reports the percentage-point change in the offered annual net bond return; the inherited return is predetermined on impact. Consumption equals output because net bonds, government debt and government purchases are zero. Fixed exposure types are initially independent of household states. Prices and transfers jointly clear the bond market and fiscal budget in the first-order equilibrium.

hold distributions. These single-innovation responses differ from the regression histories, where aggregate innovations arrive every year.

Figure C1 displays aggregate consumption and output, hours, the offered interest rate and the equal transfer. The fiscal rule links transfers proportionally to output. The return adjusts to clear the bond market as households change saving. Aggregate consumption is an equilibrium resource outcome, distinct from the exposure coefficient in Table 6. All the responses last for a really long time, because of the persistence of the aggregate shocks.

Figure C2 shows how aggregate income and consumption vary for five exposure types. The gross earnings response combines direct productivity exposure with labor adjustment.

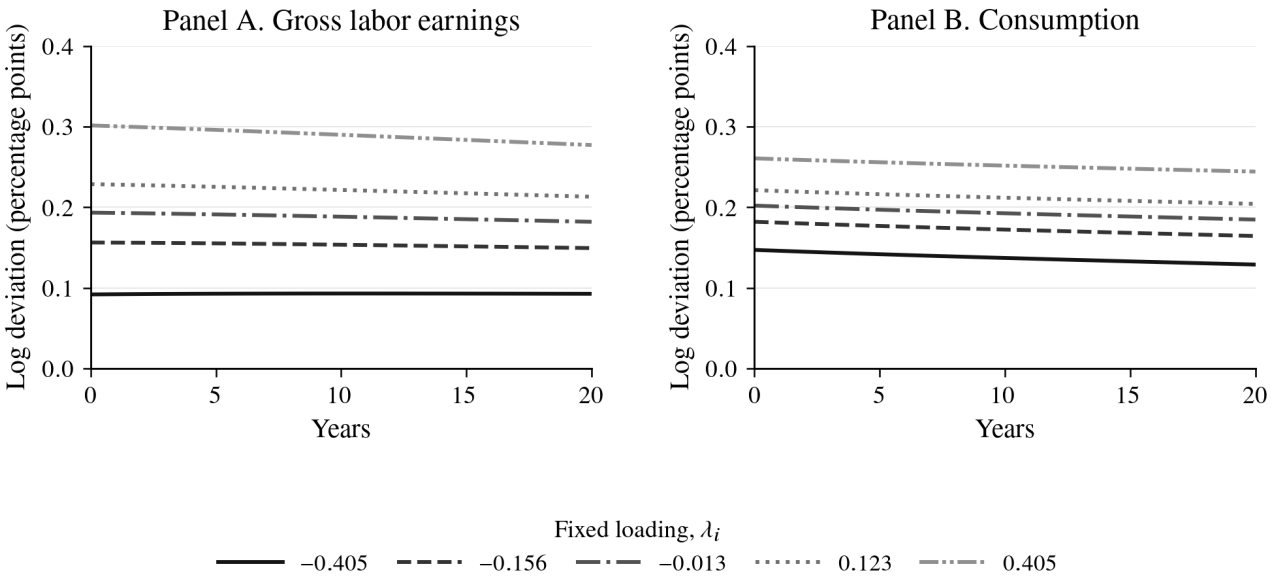


Figure C2: Earnings and consumption responses by fixed exposure

Notes: Type-conditional mean log gross earnings and log consumption responses to the innovation in Figure C1, multiplied by 100. Both panels use the same vertical scale. Legend entries are fixed loading values, not household income ranks. Each type initially has the same conditional asset and productivity distribution. Curves include endogenous hours, equilibrium interest rates, and equal transfers.

Consumption also reflects saving, taxes, and transfers. The common initial household distribution isolates exposure differences rather than confounding them with initial income or wealth sorting.