Intertemporal Self-insurance and Excess Sensitivity of Affluent Households

Ettore Savoia*

Sveriges Riksbank Working Paper Series No.464

March 12, 2026

Abstract

This paper studies household consumption behavior when saving motives differ due to preference or income risk heterogeneity. I show that self-insurance operates along two margins: a standard liquidity margin determined by the level of current resources and an intertemporal margin governed by target savings. Stronger saving motives imply higher target savings, which tilt consumption toward the future. This forward-looking behavior lowers marginal propensities to consume (MPCs) independently of current resources. Canonical incomplete-markets models confound these margins because saving motives and resources co-move. I provide conditions under standard preferences and rational expectations that overturn this mapping. When saving motives and saving capacity (such as income) are negatively related, higher-income households can be wealthier yet exhibit higher MPCs because their target savings are lower. Cross-sectional MPC–resource profiles can therefore be non-monotonic due to type sorting along the resource distribution. Empirical and quantitative applications to income-risk heterogeneity support the mechanism and highlight a new perspective in which affluent households shape aggregate demand beyond standard borrowing-constraint channels.

Keywords: Incomplete Markets, Self-Insurance, Sorting, Marginal Propensity to Consume.
JEL: D12, D81, E21

*Ettore Savoia: Research Division - Sveriges Riksbank. Postal address: Brunkebergstorg 11 - SE-103 37, Stockholm, Sweden. Contacts: ettore.savoia@riksbank.se, +46 8 787 06 98. The opinions expressed in this article are the sole responsibility of the author and should not be interpreted as reflecting the views of Sveriges Riksbank. For valuable comments and suggestions, I thank: Edouard Challe, Daria Finocchiaro, Martin B. Holm and Tullio Jappelli.

1 Introduction

A central question in macroeconomics is how household heterogeneity shapes aggregate demand. A key object for answering this question is the distribution of marginal propensities to consume (MPCs), which measure how transitory changes in household income or wealth translate into spending. In canonical heterogeneous-agent models, MPC heterogeneity is typically inferred from households' cash-on-hand (current income plus liquid savings) which summarizes their degree of self-insurance against idiosyncratic income fluctuations. Households with low levels of cash-on-hand exhibit high MPCs. This logic underlies much of the current interpretation of the joint distribution of wealth and MPCs and the way household heterogeneity works in macroeconomics.[1]

This paper revisits this interpretation by studying heterogeneity in household consumption behavior when saving motives differ due to preferences or exposure to income risk. I show that household self-insurance operates along two distinct margins. The first is the familiar liquidity margin, governed by the level of cash-on-hand. The second is an intertemporal margin governed by long-run desired savings (hereafter, target savings). Stronger saving motives imply higher target savings and a consumption path tilted toward the future. This forward-looking behavior provides greater intertemporal insurance and lowers MPCs in the Euler region independently of current cash-on-hand. As a result, beyond borrowing constraints, cross-sectional MPC–resource profiles need not reflect liquidity alone and can be non-monotonic due to the sorting of household types with different saving motives along the aggregate resource distribution.

Separating these two margins is difficult in canonical one-type models. In those models, households share common preferences and face a common income process, so precautionary motives arise only from idiosyncratic shock realizations. In that case, stronger precautionary motives are typically associated with low cash-on-hand states, especially near the borrowing constraint, where households rebuild buffers and MPCs are high. Liquidity and intertemporal margins are therefore observationally indistinguishable (Carroll, 2001). A similar difficulty arises in benchmark calibrations with ex-ante heterogeneity. In this setting as well, saving motives and resources co-move: types with stronger motives accumulate more wealth. Aggregation then mechanically sorts these types into higher wealth levels, making the cross-sectional MPC–cash-on-hand relationship appear monotone and liquidity-driven.[2]

I provide conditions under standard preferences and rational expectations that overturn this mapping. When saving motives and saving capacity (such as income) are negatively related, higher-income households can be wealthier yet exhibit higher MPCs because their target savings are lower.

The mechanism operates as follows. Households with stronger saving motives have higher target savings. When these motives are coupled with lower accumulation capacity, however, accumulated wealth need not be higher in equilibrium. Some households may

[1] See Auclert et al. (2025) for a recent survey of the literature.

[2] See, e.g., Krusell and Smith (1998); Carroll et al. (2017); Kaplan and Violante (2021); Aguiar et al. (2020); Colarieti and Monacelli (2023); Kosar and Melcangi (2024); Jappelli et al. (2024) for calibrations with ex-ante heterogeneity in preferences or income risk that preserve a monotone MPC–wealth mapping.

remain poor because they have low income and therefore cannot accumulate the wealth implied by their target savings. Nonetheless, stronger saving motives tilt consumption toward the future. This behavior generates greater intertemporal insurance and lower MPCs within the Euler region, largely independently of current resources. As a result, the cross-sectional MPC profile may be non-monotonic, or even increasing, in wealth, cash-on-hand, or income, depending on how households with different saving motives are distributed across the resource distribution.

The paper develops this mechanism in two steps. It first studies a tractable two-period model in the spirit of Leland (1968), which provides a proof of concept. The Euler equation yields a closed-form MPC and two sufficient statistics, which I term the intertemporal insurance coefficient and its local counterpart, intertemporal prudence. These objects capture how saving motives shape marginal consumption responses through the degree of intertemporal smoothing. Away from the borrowing constraint, households that tilt consumption more toward the future exhibit lower expected future curvature of marginal utility relative to current curvature. A relatively lower expected future curvature implies greater intertemporal insurance and automatically implies lower MPCs independently of current cash-on-hand. The coefficient of intertemporal *insurance* therefore serves as a sufficient statistic for MPCs in the interior region, where households smooth consumption but insurance is not yet perfect. Instead, the coefficient of intertemporal *prudence* captures the MPC differential across policies that arises from differences in front-loading at a given level of cash-on-hand. At the boundaries, the liquidity and intertemporal margins coincide. When the borrowing constraint binds, MPCs approach unity independently of saving motives. At very high levels of cash-on-hand, consumption becomes perfectly smoothed and households are effectively fully insured against income fluctuations. So MPCs approach their lower bound independently of saving motives.

The paper then turns to a standard infinite-horizon heterogeneous-agent model with ex-ante heterogeneity in preferences or labor income risk. This framework disciplines the mechanism along the cash-on-hand distribution and allows aggregation to be studied explicitly. Consistent with the pointwise intuition from the two-period framework, the model establishes the existence of equilibria in which households are wealthier due to higher income yet exhibit higher MPCs because their saving motives are weaker. A Simpson's paradox logic then arises: within each type, MPCs decline with resources, yet the cross-sectional MPC–cash-on-hand relationship can be flat or non-monotonic because resource bins are composition-weighted averages of types that sort differently along the aggregate distribution. The paper characterizes the conditions under which this non-monotonicity emerges. Identifying the two margins also implies that the usual mapping from low resources to high MPCs can be misleading. Having low resources is not equivalent to being constrained when evaluating MPCs. At the bottom of the cash-on-hand distribution, households may in fact be unconstrained and can exhibit lower MPCs than other unconstrained households with higher resources. MPCs approach unity only when borrowing constraints bind.

After establishing the general results based on preference and risk heterogeneity, the paper evaluates these predictions in a concrete application based on income-risk heterogeneity along the income distribution. This setting is natural because income risk

is a canonical and measurable driver of precautionary saving (Kimball, 1990a; Carroll and Kimball, 1996), while income is a key determinant of accumulation capacity (Dynan et al., 2004). Using Italian micro data from the Survey of Household Income and Wealth (SHIW), the paper documents new facts consistent with the proposed theoretical mechanism.

I find that income risk is heterogeneously persistent and negatively correlated with income levels, and thus with savings and cash-on-hand.[3] This correlation generates precisely the misalignment that allows identification of the two margins. Consistent with the intertemporal margin, poorer high-risk households exhibit higher expected consumption growth, suggesting that low resources are not driven by low target savings. At the same time, these low-cash, high-risk households display lower MPCs, suggesting that MPCs are in fact driven by the intertemporal margin of self-insurance. As a result, the unconditional MPC–cash-on-hand profile is flat to mildly increasing, reflecting the sorting of high-cash, low-risk households at the top percentiles. A quantitative standard incomplete-markets model with one asset and joint heterogeneity in income and income risk reproduces these aggregate MPC patterns while matching group-specific averages of liquid asset-to-income ratios and MPCs. The model also matches the observed degree of cross-sectional inequality.

The findings of this paper have clear implications for aggregate demand dynamics, although a full quantitative assessment is left to future research. The logic is straightforward. If high MPCs can arise among affluent households through heterogeneity in target savings and in the ability to reach those targets, the transmission of aggregate shocks need not operate primarily through traditional borrowing-constraint channels. Because affluent households account for a large share of aggregate consumption, their marginal responses may be quantitatively important for fiscal and monetary policy. The paper therefore offers a complementary perspective and calls for further research on the behavior of affluent households and their role in shaping aggregate demand dynamics.[4]

This perspective also helps rationalize a growing body of puzzling evidence that is difficult to reconcile with full-information rational-expectations frameworks. Many studies document that affluent households often exhibit unexpectedly high MPCs, and that a large share of MPC dispersion reflects unobserved heterogeneity rather than differences in asset holdings (Kueng, 2018; Bunn et al., 2018; Lewis et al., 2019; Fuster et al., 2021; Boutros, 2021; Colarieti et al., 2024; Boehm et al., 2025).[5] Because precautionary motives are largely unobserved, the mechanism provides a structural account of how such hetero-

[3]High income risk among low-income households is also supported by existing studies (Mueller, 2017; Broer et al., 2022; Arellano et al., 2022; Wang, 2023; Faia and Shabalina, 2023).

[4]The intertemporal insurance channel differs from the wealthy hand-to-mouth channel in (Kaplan et al., 2014, 2018). Wealthy hand-to-mouth households have low liquid resources despite holding substantial illiquid assets. In standard one-type models, including two-asset environments, high MPCs primarily arise from low levels of liquid resources. In contrast, the mechanism emphasized in this paper operates through heterogeneity in target savings. Appendix B shows that this mechanism also carries over to a two-asset environment.

[5]The evidence on the relationship between MPCs and household resources is mixed. Recent findings by Bilbiie et al. (2024) show that MPCs are low only for the top deciles of the liquid resource distribution, while they remain relatively flat across the rest of the distribution. See Crawley and Theloudis (2024) for a recent survey of this extensive literature.

geneity operates within standard models and rationalizes MPC patterns often attributed to behavioral departures.[6]

The compositional effects highlighted here have another important implication: households commonly viewed as prime targets for aggregate-demand stimulus may in fact exhibit low MPCs. Confounding precautionary motives with wealth may therefore contribute to overstated MPC estimates in the 2001 and 2008 U.S. rebates, complementing the recent critique by (Orchard et al., 2025). Consistent with this mechanism, Iao (2024) documents lower MPCs among 2008 rebate recipients facing higher unemployment risk.

Related works. This paper relates to the literature studying households' spending behavior. While much of this work focuses on liquidity-constrained households, the paper instead studies the marginal spending behavior of affluent, unconstrained households. Closely related is Aguiar et al. (2024), who emphasize curvature heterogeneity, focusing on hand-to-mouth households and preference differences. Here the emphasis is on unconstrained households and heterogeneity in income risk. In their framework, preference heterogeneity is unrelated to income levels, so more patient or risk-averse households accumulate more wealth and display lower MPCs on average. This paper shows that MPCs can be lower for more patient or more risk-averse households even when they are poorer. These results clearly contribute to the vast literature on MPC heterogeneity (see Crawley and Theloudis, 2024).

The analytical results derived in the two-period model extend Kimball (1990b) to allow for type heterogeneity. The application to income risk heterogeneity provides an example of how to separate precautionary motive from liquidity effects on the MPCs, Carroll (2001); Carroll et al. (2021) and contributes to the literature on consumption under risk (Zeldes, 989a; Kimball, 1990a,b; Carroll and Kimball, 1996). The empirical application complements empirical work linking MPCs to income-risk proxies such as self-employment (Friedman, 1957; Sala and Trivín, 2021), subjective uncertainty (Berger-Thomson et al., 2009; Jappelli et al., 2024; Kosar and Melcangi, 2024), and unemployment risk (Iao, 2024), which often—but not always—find higher risk associated with lower MPC.[7]

Overall, this paper documents empirical patterns that are difficult to reconcile with single-type precautionary-saving models but are consistent with a fully rational framework once income and saving motives vary jointly. The results therefore speak directly to the burgeoning heterogeneous-agent literature (see Auclert et al. (2025)) and highlight the need for further research on the role of affluent households in aggregate dynamics. This perspective is particularly relevant in light of recent empirical evidence from administrative data pointing to the important role of high-income households in shaping

[6]This does not diminish the importance of behavioral mechanisms, among others, such as present bias (Laibson et al., 2024), inattention (Auclert et al., 2020), or other consumption anomalies (Andre et al., 2025).

[7]Online Appendix of this paper documents cases in which MPC–risk monotonicity can fail—for instance when an absorbing state (e.g. unemployment) shifts mass toward high-MPC states—potentially generating a U-shaped MPC–risk pattern. The results in this paper suggest that other non-monotonic MPC–risk patterns may reflect richer income dynamics that go beyond the scope of this paper but represent an interesting avenue for future research, such as hidden covariates or nonlinearities in income risk (see De Nardi et al., 2020; Caplin et al., 2023; Gálvez and Paz-Pardo, 2023 for evidence on income dynamics). Other channels, such as idiosyncratic expenditure innovations (Bergholt et al., 2025), may also play an important role and deserve further investigation.

aggregate fluctuations (Amberg et al., 2022; Andersen et al., 2023).

The remainder of the paper is organized as follows. Section 2 introduces the two-period model, and Section 3 the infinite-horizon model. Sections 4 and 5 provide the empirical and quantitative validation of the mechanism, respectively. Section 6 concludes.

2 An illustrative two-period model

This section derives closed-form sufficient statistics that formalize the central role of intertemporal insurance in driving MPC heterogeneity across types. For analytical transparency, the analysis is confined to a two-period setting in the spirit of Leland (1968) and Sandmo (1970), while all numerical and quantitative results rely on the infinite-horizon framework, where the same mechanism carries through. The objective is to isolate the drivers of the MPC by comparing types operating under different policy functions and at different wealth levels.

Environment. Households live for two periods, choose current and next-period consumption $\{c_0, c_1\}$ and savings a_1, and solve

$$\max_{\{c_0,c_1,a_1\}} \quad u(c_0) + \beta\, \mathbb{E}_0[u(c_1)] \tag{1}$$

$$\text{s.t.} \quad c_0 + a_1 = x_0, \tag{2}$$

$$c_1 = Ra_1 + y_1, \tag{3}$$

$$a_1 \geq 0. \tag{4}$$

Here, c_0 denotes current consumption, and x_0 is cash-on-hand in the current period, given by the sum of deterministic income y_0 and beginning-of-period assets a_0. The agent allocates x_0 between consumption c_0 and savings a_1. Next-period consumption c_1 is financed by capitalized savings Ra_1, where $R = 1 + r$ is the gross interest rate, and by stochastic income y_1 received in the next period. The expectation $\mathbb{E}_0$ is taken with respect to the distribution of future income conditional on the information available in period 0. Unless otherwise specified, I assume $\beta R < 1$. While this condition is not needed for the two-period problem, it is maintained for consistency with the infinite-horizon environment studied later in the paper. Whenever the borrowing constraint is slack, optimality is characterized by the Euler equation

$$u'(c_0) = \beta R\, \mathbb{E}_0[u'(c_1)]\,. \tag{5}$$

Here $u'' < 0$ and $u'(x_0 - a_1)$ is strictly increasing in a_1, while $\beta R\, \mathbb{E}_0[u'(Ra_1 + y_1)]$ is strictly decreasing in a_1. Hence, the optimal savings choice a_1^* is uniquely determined. The policy function is therefore given by:

$$c_0 = x_0 - a_1^*. \tag{6}$$

Drivers of saving motives. Throughout this paragraph, I work in general terms using Equation (5). It is nevertheless useful to make it clear which forces drive saving motives. First, note that any parameter change that raises the right-hand side relative to the left-hand side at a given allocation induces higher desired saving (higher target savings), thereby lowering c_0 and shifting the consumption rule toward less front-loading.

With CRRA utility $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$ where $1/\gamma$ is the intertemporal elasticity of substitution (IES), equation (5) becomes $1 = \beta R\, \mathbb{E}_0[(\frac{c_1}{c_0})^{-\gamma}]$. By taking logs of the latter and applying the standard lognormal approximation one obtains

$$\mathbb{E}_0 \Delta \log c_1 \;\approx\; \frac{1}{\gamma} \log(\beta R) \;+\; \frac{\gamma}{2}\, \mathbb{V}_0(\Delta \log c_1),$$

where $\Delta \log c_1 \equiv \log c_1 - \log c_0$. The latter shows that stronger saving motives (type heterogeneity) can arise from (i) greater patience (higher β), (ii) higher risk aversion or lower IES (higher γ when $\beta R < 1$), and (iii) greater consumption risk (higher $\mathbb{V}_0(\Delta \log c_1)$), which is the reduced-form object through which higher income risk raises precautionary saving.

MPC derivation. Differentiating 6 with respect to x_0 gives a general expression for the current-period MPC:

$$\frac{dc_0}{dx_0} = 1 - \frac{da_1^*}{dx_0}.$$

Differentiating the Euler equation (5) with respect to x_0, one recovers an expression for the marginal propensity to save (MPS) $\frac{da_1^*}{dx_0}$. Hence, after some straightforward algebra, the current-period MPC is

$$MPC \equiv \frac{dc_0}{dx_0} = \frac{\beta R^2 \mathbb{E}_0[u''(c_1)]}{u''(c_0) + \beta R^2 \mathbb{E}_0[u''(c_1)]} \in (0, 1). \tag{7}$$

where $u''(c_0)$ is the second derivative (local curvature) of utility with respect to current consumption, and $E_0[u''(c_1)]$ captures the expected second derivative of future utility under uncertainty. Under $u''(\cdot) < 0$ and standard regularity conditions (in particular, $\mathbb{E}_0[|u''(c_1)|] < \infty$), both the numerator and denominator in (7) are negative and the denominator exceeds the numerator in absolute value, so the MPC lies strictly between 0 and 1. Since the formula is derived from the Euler equation, it applies only to unconstrained households. When the constraint binds ($a_1 = 0$, hence $c_0 = x_0$), the MPC is exactly 1 by definition. If expected income is increasing (so that $\mathbb{E}_0[c_1] > c_0$ in the constrained region), consumption growth is positive there. Beyond the constraint, the Euler equation governs intertemporal allocation through expected marginal utility rather than expected consumption levels. Along the consumption policy of a given type, states with more front-loaded consumption (i.e. lower expected consumption growth) are associated with a lower MPC. However, this relationship is local and within-type: it describes movements along a given type's policy function and need not hold when comparing households

of different types, as considered below.[8]

MPC heterogeneity across *types*. Equation (7) characterizes the MPC of unconstrained households from the Euler equation. I use it to compare MPCs across households of different *types*, abstracting from the underlying source of stronger saving motives, as well as from whether and how these forces covary with income levels. For now, the goal is to identify the drivers of MPC differences across types facing different policy rules.

Proposition 1 (Intertemporal Insurance as a sufficient statistic for global cross-type MPC ranking)**.** *Consider two types A and B in the interior region of cash-on-hand $\mathcal{X} \subset \mathbb{R}_{++}$, so that for both types the MPC is given by (7). Define the* coefficient of intertemporal insurance *as*

$$\mathcal{I}_i^{\mathcal{I}} \equiv \frac{-u''(c_0^i)}{\beta_i R_i^2 \mathbb{E}_0[-u''(c_1^i)]} > 0, \qquad \Rightarrow \qquad MPC_i = \frac{1}{1+\mathcal{I}_i^{\mathcal{I}}}. \tag{8}$$

Then, for any pair $\{x_{0,A}, x_{0,B}\} \in \mathcal{X}$, we have:

$$MPC_B(x_{0,B}) > MPC_A(x_{0,A}) \iff \mathcal{I}_B^{\mathcal{I}}(x_{0,B}) < \mathcal{I}_A^{\mathcal{I}}(x_{0,A}). \tag{9}$$

The coefficient $\mathcal{I}^{\mathcal{I}}$ summarizes how much intertemporal smoothing a household achieves at its current state. With common preferences and interest rates, it simply says that households facing relatively low expected future curvature $E_0[-u''(c_1)]$ display lower MPCs today. The logic is that stronger saving motives raise desired savings and reduce front-loading, implying a reduction in c_0. This behavior lowers expected future curvature relative to current curvature, reflecting stronger intertemporal insurance (i.e. $\mathcal{I}^{\mathcal{I}}$ increases) and implies a lower MPC regardless of cash-on-hand.

However, when comparing across policy functions, this object also reflects the within-type effect that MPC declines with cash-on-hand due to the concavity of consumption rules. Under standard curvature (e.g., $u'' < 0$ and $u''' > 0$), both $|u''(c_0)|$ and $|\mathbb{E}_0[u''(c_1)]|$ decline with cash-on-hand. Because $\mathcal{I}^{\mathcal{I}}$ depends on *both* terms, it can be that precautionary and liquidity forces may dominate or offset in different regions: households may exhibit similar intertemporal insurance and thus similar MPCs at different cash levels, while households at the same cash level can display different MPCs if they differ in intertemporal insurance.

Importantly, because objects in this section capture a pointwise comparison, results are conceptual: the coefficient of self-insurance clarifies which margins can shift MPCs across types operating under different consumption rules, but it cannot by itself speak to aggregate MPC patterns or imply a uniform MPC ranking at a given level of cash-on-hand across policy functions. Hence, the key information to extract is that, across types, higher assets do not automatically imply stronger intertemporal insurance and therefore lower MPCs. The relevant question now is where curvature dominates and where liquidity does. To this end, it is useful to compare types *at a given cash level*, to isolate differences in the consumption rule itself.

[8] See also Aguiar et al. (2024) for a more detailed discussion on this.

Proposition 2 (Intertemporal prudence and local MPC ranking)**.** *Consider two types A and B in the interior (Euler-equation) region $\mathcal{X} \subset \mathbb{R}_{++}$. Fix a level of cash-on-hand $\bar{x}_0 \in \mathcal{X}$.*

Define the local intertemporal prudence coefficient

$$\mathcal{P}_i^{\mathcal{I}}(\bar{x}_0) := -\frac{u_i'''(c_0^i)}{u_i''(c_0^i)}\Big(c_0^i - \mathbb{E}_0[c_1^i]\Big), \qquad i \in \{A, B\}.$$

Under a first-order Taylor expansion of expected future curvature $\mathbb{E}_0[-u_i''(c_1^i)]$ around c_0^i, intertemporal insurance satisfies

$$\mathcal{I}_i^{\mathcal{I}} \approx \frac{1}{\beta_i R_i^2\big(1 + \mathcal{P}_i^{\mathcal{I}}\big)}.$$

Therefore,

$$MPC_B(\bar{x}_0) > MPC_A(\bar{x}_0) \quad \Longleftrightarrow \quad \beta_B R_B^2\big(1 + \mathcal{P}_B^{\mathcal{I}}\big) > \beta_A R_A^2\big(1 + \mathcal{P}_A^{\mathcal{I}}\big).$$

In particular, if $\beta_A R_A^2 = \beta_B R_B^2$, then

$$MPC_B(\bar{x}_0) > MPC_A(\bar{x}_0) \quad \Longleftrightarrow \quad \mathcal{P}_B^{\mathcal{I}}(\bar{x}_0) > \mathcal{P}_A^{\mathcal{I}}(\bar{x}_0).$$

Thus, holding cash fixed, a higher $\mathcal{P}^{\mathcal{I}}$ means lower intertemporal insurance, $\mathcal{I}^{\mathcal{I}}$, and therefore a higher MPC. This makes it clear that MPCs are independent of precautionary motives at both extremes: when the borrowing constraint binds and under perfect smoothing, that is, when current and future curvature coincide. In the interior, however, MPCs are governed by intertemporal consumption behavior and therefore by the implied ratio of current curvature to expected future curvature. Interestingly, in the limit in which consumption smoothing equalizes consumption across periods, households converge to full intertemporal insurance: $\mathcal{P}^{\mathcal{I}} \to 0$ and $\mathcal{I}^{\mathcal{I}} \to \frac{1}{\beta R^2}$, while consumption growth satisfies $\mathbb{E}_0[c_1 - c_0] \to 0$. The MPC therefore converges to $\frac{\beta R^2}{1+\beta R^2}$, which equals $1/2$ when $\beta R^2 = 1$. This means that, in the limit, the allocation under stronger precautionary motives converges to the certainty-equivalent allocation, whereby consumers are effectively fully insured across periods.

Analogy with the concept of Absolute Prudence. To a first order, the analogy between the coefficient of intertemporal prudence $\mathcal{P}^{\mathcal{I}}$ and the coefficient of absolute prudence introduced by Kimball (1990b), $P_A(c) \equiv -u'''(c)/u''(c)$, is immediate. Decreasing absolute prudence ($P_A(c)$) governs within-type precautionary behavior across cash levels holding curvature fixed; intertemporal prudence ($\mathcal{P}^{\mathcal{I}}$) governs between-type precautionary behavior across different curvatures holding cash fixed. Within a type, households with greater absolute prudence ($P_A(c)$) hold fewer current resources, have lower static self-insurance, and therefore exhibit higher MPCs. Across types, households with greater intertemporal prudence ($\mathcal{P}^{\mathcal{I}}$) have lower target savings, less intertemporal insurance, and

thus higher MPCs.[9]

CRRA closed-form expressions. Under CRRA preferences, $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$, both coefficients admit closed-form expressions:

$$\mathcal{I}^{\mathcal{I}} = \frac{c_0^{-\gamma-1}}{\beta R^2 \, \mathbb{E}_0[c_1^{-\gamma-1}]}, \qquad \mathcal{P}^{\mathcal{I}} = \frac{\gamma+1}{c_0} \left(c_0 - \mathbb{E}_0[c_1]\right).$$

Implications. These results have clear empirical and quantitative implications. Empirically, the intertemporal insurance coefficient ($\mathcal{I}^{\mathcal{I}}$) implies that reduced-form MPC–wealth profiles may reflect compositional effects: a flat or increasing relationship can arise when types with lower interior MPCs are concentrated at lower resources, while wealthier households display higher MPCs because they are less intertemporally insured. Intertemporal prudence ($\mathcal{P}^{\mathcal{I}}$) instead provides a disciplined interpretation of residual MPC heterogeneity in fixed-cash designs. In quantitative models with ex-ante heterogeneity in saving motives, aggregate MPCs depend jointly on policy rules and the stationary distribution. When saving motives and accumulation capacity are misaligned, households with lower wealth may still display lower MPCs due to stronger intertemporal insurance, implying that the aggregate MPC–wealth relationship ultimately reflects the sorting of types along the wealth distribution. The next sections illustrate these mechanisms with practical applications.

3 The incidence of saving motives and its implications for the MPC

This section evaluates the insights of the two-period model but accounts for distributional dynamics using an infinite-horizon standard incomplete-markets framework (SIM hereafter). The environment corresponds to the household block of canonical heterogeneous-agent (HA) frameworks, both in the classical incomplete-markets tradition of Huggett (1993) and Aiyagari (1994), and in New Keynesian (HANK) models Auclert et al. (2021, 2025). The analysis is conducted in partial equilibrium. The discussion proceeds through comparative statics to isolate the underlying mechanisms and to motivate the empirical analysis that follows. The quantitative analysis based on the data is introduced in Section 5.

Environment. Each type defined below is populated by a unit measure of infinitely lived households $i \in [0,1]$ facing idiosyncratic income risk and trading a single non-state-

[9] Note that to a first order, $\mathcal{P}^{\mathcal{I}} \approx P_A(c_0)\,(c_0 - \mathbb{E}_0[c_1])$. Hence it is not a monotone transform of absolute prudence: it depends on preferences through $P_A(c_0)$ and on the equilibrium allocation through front-loading (and, more generally, through $\mathbb{E}_0[-u''(c_1)]$). Hence changes in $\mathcal{P}^{\mathcal{I}}$ need not map one-to-one into changes in P_A. Like for the domain, the approximation requires $1 + \mathcal{P}_i^{\mathcal{I}} > 0$. Under CRRA preferences, $\mathcal{P}_i^{\mathcal{I}} \geq 0$ whenever consumption is front-loaded in expectation ($c_0 \geq \mathbb{E}_0[c_1]$). If instead consumption is back-loaded in expectation, then $\mathcal{P}_i^{\mathcal{I}} < 0$, and the first-order approximation remains well behaved only if $\mathcal{P}_i^{\mathcal{I}} > -1$. Cases with $\mathcal{P}_i^{\mathcal{I}} \leq -1$ should be interpreted as regions in which higher-order terms of the Taylor expansion are non-negligible.

contingent asset. The household's problem is given by:

$$V_t(y_{i,t}, a_{i,t-1}) = \max_{c_{i,t}, a_{i,t}} u(c_{i,t}) + \beta E_t[V_{t+1}(y_{i,t+1}, a_{i,t})] \tag{10}$$

subject to:

$$c_{i,t} + a_{i,t} = (1 + r)a_{i,t-1} + \psi \cdot y_{i,t} \tag{11}$$

$$a_{i,t} \geq 0. \tag{12}$$

Consumption is denoted by $c_{i,t}$, stochastic labor income by $y_{i,t}$, lagged assets by $a_{i,t-1}$, and current savings by $a_{i,t}$. The interest rate r is exogenous. Preferences are represented by an isoelastic utility function, $u(c_{i,t}) = \frac{c_{i,t}^{1-\gamma}}{1-\gamma}$, where β denotes the discount factor and the intertemporal elasticity of substitution (IES) equals $1/\gamma$. The model's stationary distribution is computed following Auclert et al. (2021). Without loss of generality, households cannot borrow and the stochastic income process follows a logarithmic AR(1) process, discretized using the Rouwenhorst (1995) method over seven states s:

$$logy_{i,t} = \rho \log y_{i,t-1} + \varepsilon_{i,t}, \tag{13}$$

where $\varepsilon_{i,t}$ is an i.i.d. normal process with mean zero and standard deviation σ_ε.

MPC. For each type, the MPC is computed as the local derivative of the stationary consumption policy with respect to individual cash-on-hand, holding the income state fixed, and thus holding expectations about future income constant.

Current and Permanent income. The novelty is the parameter ψ, which captures differences in the level of permanent income, or more generally in households' accumulation capacity. The parameter ψ therefore scales the level of income in the budget constraint and captures differences in permanent income across types. Throughout the paper, all moments are expressed relative to average current-income, given by the stochastic income component $y_{i,t}$ which is normalized to have unit mean, $E[y] = 1$. Under this normalization, the model-implied wealth-to-income ratio coincides with average asset holdings. Consequently, ψ affects equilibrium asset holdings and therefore the model-implied wealth-to-income ratios.[10]

Target savings and saving motives. In this setting, target savings denotes the type-specific long-run asset (steady-state) position implied by saving motives given preferences, the income process, and the interest rate. Also in this section, stronger saving motives (type heterogeneity) can arise from (i) greater patience (higher β), (ii) higher risk aversion or lower IES (higher γ when $\beta R < 1$), and (iii) greater income risk (σ).

[10]This provides the simplest possible way to model frictions in savings accumulation. Households may have a stronger desire to save, but their ability to accumulate savings is constrained (enhanced) when ψ is low (high). More complex mechanisms, such as committed consumption, documented empirically by Shore and Sinai (2010), are also possible and represent an avenue for future research.

3.1 Distributional Characterization of MPC Rankings

Conceptual map and dimensions of heterogeneity. The goal is to provide objects that capture cross-type variation and help interpret empirical estimates in the presence of unobserved heterogeneity. Consequently, I study two complementary dimensions of heterogeneity:

(1) Between-economy heterogeneity (comparative statics across stationary equilibria). First, I study how the *average* MPC and *average* resources co-move across stationary equilibria indexed by a saving-motive shifter $\tau = \tau(\beta, \gamma, \sigma)$ and an accumulation-capacity shifter ψ. This step is necessary to establish that non-monotonicity is feasible once saving motives and accumulation capacity are misaligned.

(2) Within-economy heterogeneity (types coexisting in a single stationary equilibrium). Second, I translate the between-economy insight into a within-economy characterization. The relevant objects in this case are the stationary distributions of each type and the composition of types along the cross-sectional distribution of resources.

3.1.1 Between-Equilibrium Comparative Statics

To clarify the underlying mechanisms, Figures 1–3 consider three cases: households differing only in saving motives (τ), only in permanent income (ψ), and a negative incidence case in which stronger saving motives (higher τ) are associated with lower permanent income (lower ψ).[11]

Heterogeneity in saving motives only. The top row in Figure 1 varies idiosyncratic income risk, the middle row the discount factor, and the bottom row the relative risk aversion. A common pattern emerges across all scenarios. Stronger saving motives imply a higher target savings, leading to higher expected consumption growth and a lower fraction of households at the borrowing constraint. In the absence of frictions, this translates into higher asset accumulation, more cash-on-hand, and lower average MPCs. Hence, the model does not deliver qualitatively new insights relative to the single-type case, because the intertemporal and liquidity margins of self-insurance remain confounded.[12]

Heterogeneity in the level of permanent income only. Before turning to the case of negative incidence, it is useful to clarify the role of ψ. As shown by Straub (2019), in incomplete-markets models with homothetic preferences and idiosyncratic risk, proportional differences in permanent income leave the cross-sectional MPC distribution unchanged. In fact, the consumption function satisfies

$$c(a, y, \psi) = \psi\, c\left(\frac{a}{\psi}, y, 1\right),$$

[11] The case with positive covariance is redundant, as it produces patterns similar to the saving-motive benchmark; for completeness, it is reported in Appendix B.

[12] For a detailed discussion on this issue see also Carroll (2001).

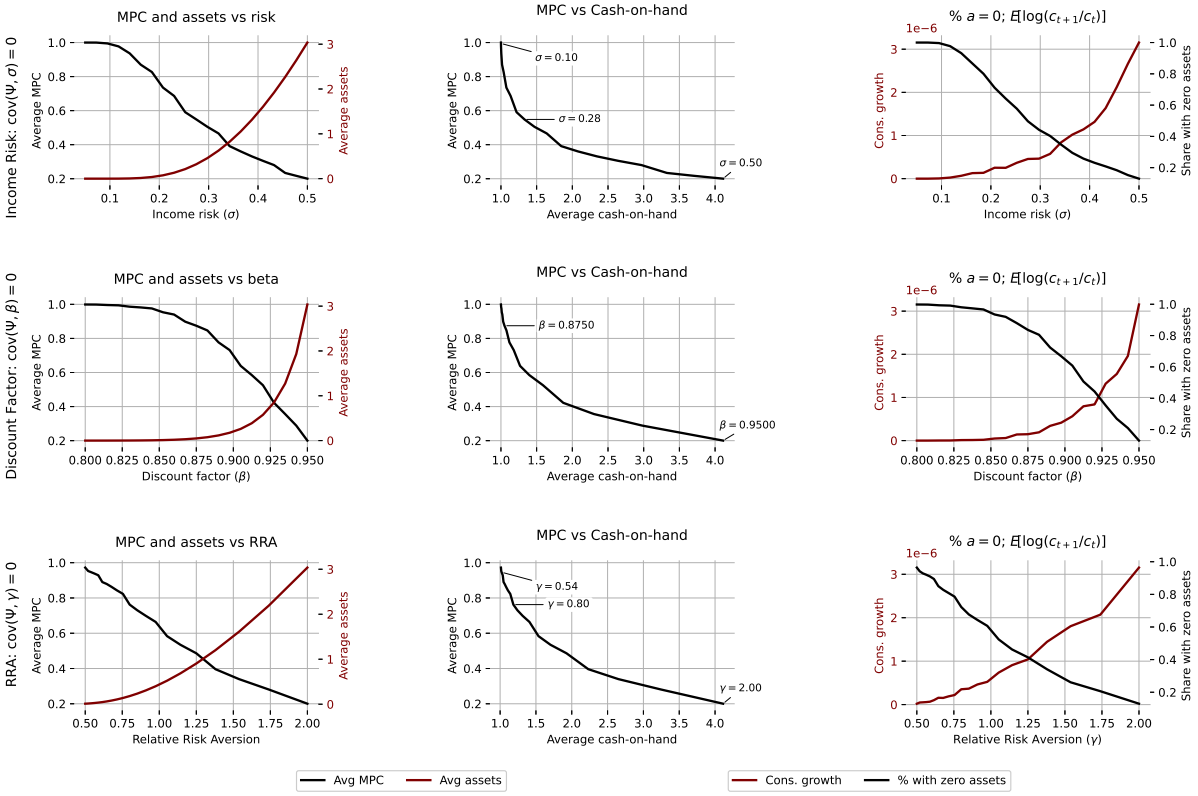


Figure 1: Consumption and Saving behavior with Heterogeneity in saving motives.

Notes: The figure reports Steady-state moments under alternative sources of heterogeneity in saving motives. Each row corresponds to a different parameter: income risk σ (top panels), the discount factor β (middle panels), and the relative risk aversion γ (bottom panels). The left column reports steady-state average MPCs (black) out of cash-on-hand, together with average assets (red), as the relevant parameter varies. The middle column shows the implied MPC–cash-on-hand schedule. The right column plots consumption growth, $\mathbb{E}\left[\log\left(c_{t+1}/c_t\right)\right]$ (red), together with the share of households with zero assets in the ergodic wealth distribution (black), as a function of the corresponding parameter. Across all exercises, the interest rate is set to $r = 2.5\%$ and the permanent-income scalar to $\psi = 1$. When income risk varies, I fix relative risk aversion at $\gamma = 1$, the persistence of the income process at $\rho = 0.95$, and the discount factor at $\beta = 0.95$. When varying the discount factor β, all other parameters remain unchanged and the standard deviation of the income process is set to 0.50. Similarly, when varying the RRA (γ), all remaining parameters are held fixed at their baseline values.

so that the stationary distribution of a/ψ and the MPC conditional on the normalized state variables $(a/\psi, y)$ are invariant to ψ. Figure 2 illustrates this homotheticity result: higher ψ scales consumption and wealth proportionally, while stationary MPC moments remain essentially unchanged. This homotheticity result applies to the benchmark case in which types differ only in income scale. The exercises below instead study joint heterogeneity in income scale and saving motives. The relevant mechanism is therefore not a pure scale effect of ψ, but a misalignment between stronger precautionary motives and lower realized asset positions: stronger precautionary motives raise desired long-run savings, while lower ψ reduces equilibrium asset holdings. Importantly, the analysis that follows concerns a gap between desired and realized asset positions, rather than a breakdown of homotheticity.

Negative incidence. I now turn to the key step in the paper: isolating the intertemporal (target-saving) margin of self-insurance in the infinite-horizon model. Proposition 3 formalizes this argument and the graphical proof follows below.

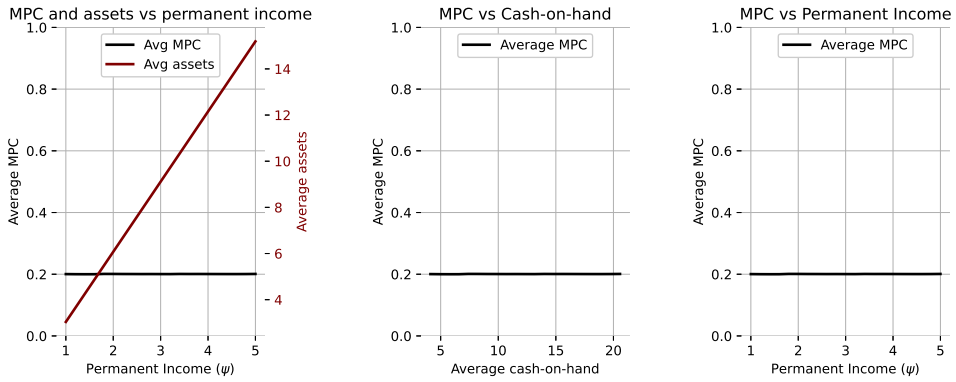


Figure 2: Consumption and Saving behavior with Heterogeneity in permanent income only.

Notes: Panels report steady-state moments as a function of the permanent income scalar ψ. The left panel shows the steady-state average MPCs (black) out of cash-on-hand, and average assets (red) across different ψ levels. The center panel reports the implied MPC–cash-on-hand relationship. The third panel shows the implied MPC–permanent-income profile.

Proposition 3 (Higher MPC in More Affluent Economies)**.** *Consider the stationary incomplete-markets economy in* (10)–(13) *with borrowing constraint* $a \geq 0$*,* $R > 1$*, and CRRA preferences. Fix all primitives except a saving-motive shifter* $\tau = \tau(\beta, \gamma, \sigma)$ *and a permanent-income shifter* $\psi > 0$*.*

Let $\bar{M}(\tau, \psi)$ *denote the stationary average MPC and* $\bar{A}(\tau, \psi)$ *average assets in the associated stationary equilibrium. Then there exist admissible pairs* (τ_H, ψ_L) *and* (τ_L, ψ_H) *such that*

$$\bar{A}(\tau_L, \psi_H) \geq \bar{A}(\tau_H, \psi_L) \quad \textit{and} \quad \bar{M}(\tau_L, \psi_H) > \bar{M}(\tau_H, \psi_L).$$

Proposition 3 argues that economies with higher permanent income but weaker saving motives, and thus lower target savings, can display higher MPCs. The intuition is that stronger precautionary motives raise target savings but need not translate into higher wealth when accumulation capacity is limited. This negative incidence separates intertemporal insurance from liquidity. If MPCs in the interior are driven by target savings, allocations can arise in which households with higher resources display higher MPCs because their target savings are lower.

Figure 3 illustrates this result. The panels report the underlying comparative statics across three sources of heterogeneity in saving motives—income risk σ, patience β, and relative risk aversion γ. In each case, stronger saving motives increase the desired degree of intertemporal self-insurance (higher target savings) and lead households to front-load consumption less. Consistent with this mechanism, expected consumption growth rises and the stationary share of households at the borrowing constraint declines (right column), even though average asset holdings need not increase one-for-one.

In fact, the left panels show that asset accumulation is hump-shaped: as saving motives strengthen, the ability to accumulate wealth declines because permanent income falls. Nevertheless, MPCs continue to decrease and approach one only in the limiting case in which accumulation capacity becomes arbitrarily small. As $\psi \downarrow 0$ (approximated in the simulations by very small positive values of ψ), household resources approach zero and the economy moves toward the high-MPC region. This highlights that being *at* the

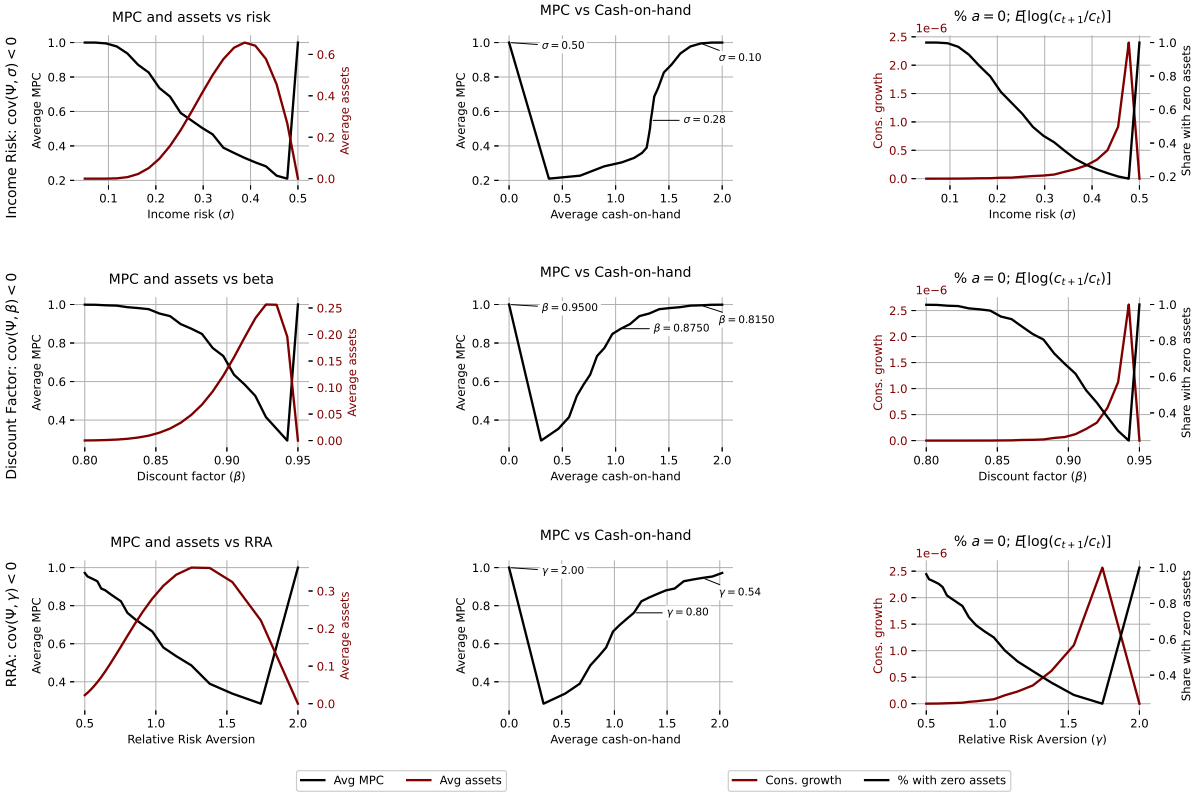


Figure 3: Consumption and Saving behavior with negative incidence of saving motives.

Notes: The figure reports steady-state moments under alternative sources of heterogeneity in saving motives combined with a negatively related permanent-income scalar ψ. Each row corresponds to a different parameter configuration: as income risk σ increases, ψ decreases from 1 to 0 (top panels); the same negative covariance is imposed when varying the discount factor β (middle panels) and the relative risk aversion γ (bottom panels). The left column reports steady-state average MPCs (black) out of cash-on-hand, together with average assets (red), as the relevant parameter varies. The middle column shows the implied MPC–cash-on-hand schedule. The right column plots expected consumption growth, $\mathbb{E}\left[\log\left(c_{t+1}/c_t\right)\right]$ (red), together with the share of households with zero assets in the ergodic wealth distribution (black), as a function of the corresponding parameter. Across all exercises, remaining parameters are held fixed at their baseline values as described under Figure 1.

borrowing constraint is not the same as being merely *close to* it.

The central panels further illustrate that cash-on-hand is not a sufficient statistic but a compositional object: the relationship between MPCs and average cash-on-hand is not governed by liquidity alone and can be flat, non-monotonic, or even increasing (as in this numerical example), depending on how high-MPC types are distributed along the resource distribution.

Robustness and additional results. Appendix B reports robustness to alternative specifications of the covariance between income levels (accumulation capacity) and saving motives, as well as to richer income processes. The results also hold when income levels are negatively correlated with either persistent or transitory volatility, and extend to standard model enrichments such as endogenous labor supply, heterogeneous borrowing limits, a two-asset structure, and unemployment risk. Similar comparative statics arise under alternative forms of limited accumulation capacity. For instance, the same patterns obtain when, instead of varying the permanent-income parameter, the ergodic mean of current income declines as saving motives increase, supporting the more general insight

that when stronger target-saving motives are not matched by a proportional capacity to accumulate buffers, the two margins of self-insurance can be identified separately.

3.1.2 Within-Equilibrium Results

Proposition 3 establishes an across-equilibrium existence result for non-monotonicity between MPCs and resources. The relevant question, however, is whether the same misalignment logic operates within a given economy, and under what conditions type-level differences generate non-monotonicity in the aggregate MPC–wealth profile.

Proposition 4 (Type ranking in stationary averages within equilibrium)**.** *Fix a stationary equilibrium with a finite set of types $i \in \mathcal{T}$. For each type i, let $s_t \equiv (a_{t-1}, \xi_t)$ denote the individual state, where a_{t-1} is beginning-of-period assets and ξ_t collects the idiosyncratic income state. Let $m_i(s_t)$ denote the marginal propensity to consume implied by the policy function at state s_t, and let F_i denote the stationary distribution of states.*

Define the stationary mass at the borrowing constraint $a_{\min}$ as

$$D_{0,i} \equiv \Pr_i(a_t = a_{\min}),$$

and the average MPC in the interior (Euler) region as

$$\bar{m}_i^E \equiv \mathbb{E}_i[m_i(s_t) \mid a_t > a_{\min}].$$

When the borrowing constraint binds ($a_t = a_{\min}$), consumption adjusts one-for-one with cash-on-hand, implying $m_i(s_t) = 1$ at the constraint. The type-i average MPC therefore admits the decomposition

$$\bar{m}_i \equiv \mathbb{E}_i[m_i(s_t)] = D_{0,i} + (1 - D_{0,i})\,\bar{m}_i^E.$$

Now consider two types i and j. If type i has stronger intertemporal insurance in the Euler region, in the sense that

$$\bar{m}_i^E < \bar{m}_j^E,$$

then $\bar{m}_i < \bar{m}_j$ if and only if

$$(D_{0,j} - D_{0,i}) > (1 - D_{0,i})\,\bar{m}_i^E - (1 - D_{0,j})\,\bar{m}_j^E.$$

Interpretation. Proposition 4 provides a type-level decomposition of average MPCs into an extensive margin (the stationary mass at the borrowing constraint) and an intensive margin (the average MPC in the Euler region). The decomposition yields a sufficient condition under which a type can be relatively poor yet still display a low average MPC if stronger saving motives sufficiently reduce its interior MPC and/or its stationary mass at the borrowing constraint.

Remarks. First, because within-type MPCs decline with resources and are independent of target savings at the tails of the resource distribution, type-specific MPC–wealth

schedules may cross locally near the borrowing constraint and converge at the top of the distribution, where consumption functions become linear.[13] Such local comparisons along policy rules are therefore not informative about cross-type MPC rankings. Second, non-monotonicity in cross-type averages does not necessarily imply a non-monotonic aggregate MPC–wealth gradient. The proposition concerns type-level stationary averages and does not pin down the slope of the aggregate MPC–cash profile, which depends on the sorting of types across cash bins, as formalized below.

Proposition 5 (Within-equilibrium aggregation and sorting)**.** *Consider a stationary equilibrium with a finite or countable set of types $i \in \mathcal{T}$, type-specific MPC schedules $m_i(a)$, and type shares conditional on assets $\pi_i(a) \equiv \Pr(i \mid a)$. Thus, $m_i(a)$ denotes the MPC of type i at asset level a, while $\pi_i(a)$ denotes the conditional population share of type i among households with assets a. Define the aggregate MPC schedule at asset level a as*

$$M(a) \equiv \mathbb{E}[m \mid a] = \sum_{i \in \mathcal{T}} \pi_i(a)\, m_i(a).$$

If M is differentiable, then

$$M'(a) = \underbrace{\sum_{i \in \mathcal{T}} \pi_i(a)\, m_i'(a)}_{\textit{within-type effect}} + \underbrace{\sum_{i \in \mathcal{T}} \pi_i'(a)\, m_i(a)}_{\textit{sorting effect}}.$$

Hence, the aggregate MPC schedule is locally increasing in a (i.e., $M'(a) > 0$) if and only if the sorting effect dominates the average within-type decline, that is,

$$\mathrm{Cov}_a(m_i(a),\, \partial_a \log \pi_i(a)) > -\mathbb{E}_a[m_i'(a)]\,.$$

It is locally flat if equality holds and locally decreasing otherwise.

Above, the terms $\mathbb{E}_a[\cdot]$ and $\mathrm{Cov}_a(\cdot,\cdot)$ denote the expectation and covariance taken across types conditional on asset level a.

Interpretation. The proposition characterizes the condition under which the aggregate MPC–asset profile is increasing, flat, or decreasing within a stationary equilibrium. The first term, the within-type effect, captures the mechanical decline of MPCs with assets within each type. Absent compositional changes, this term alone implies a decreasing aggregate MPC schedule. The second term, the sorting effect, captures how the conditional composition of types varies across asset levels in the cross-sectional asset distribution. If high-MPC types become relatively more prevalent at higher asset levels, the sorting term is positive and can offset or overturn the within-type decline. Hence, non-monotonicity in the aggregate MPC–wealth profile is not a property of individual policy functions but of the joint distribution of types across the resource distribution. The aggregate gradient is increasing if and only if the covariance between type-specific MPCs and the relative change in their population shares along the asset distribution is sufficiently positive to

[13] I thank Edouard Challe for helpful discussions on this point.

dominate the average within-type decline in MPCs, implying a Simpson's paradox (see Appendix B.9.1).

Remarks. Proposition 5 is purely distributional. It does not depend on the number of types or on whether types are fixed or change over time. It is also independent of whether the environment is partial or general equilibrium: equilibrium feedback may affect the stationary distribution through prices, but once the joint distribution of types across assets is given, the condition for non-monotonicity remains algebraic. Its empirical relevance depends on whether plausible primitives generate compositional shifts across the asset distribution that affect the slope of the aggregate MPC profile. The proposition therefore provides a structural diagnostic for when the slope can change sign.

Taking stock. Propositions 4 and 5 are stated in general terms, but in the present setting their economic content becomes relevant when the monotonic mapping from saving motives to realized wealth is broken, as in Proposition 3. If households with stronger saving motives also accumulate proportionally more wealth, heterogeneity is absorbed by liquidity and the standard monotone MPC–wealth mapping obtains. The possibility of poor low-MPC types, and thus of a non-monotonic aggregate MPC–wealth profile, arises when stronger saving motives raise target savings without generating proportionally higher wealth. The empirical and quantitative relevance of this mechanism must therefore be assessed in concrete applications, as done in the remainder of the paper.

Relation to the classical buffer–stock framework. The mechanism studied so far is related to the classical buffer–stock formulation of Carroll (1997), but the present setting makes the interaction between saving motives and accumulation capacity more transparent. In the normalized buffer–stock model, wealth is expressed relative to permanent income ψ_t, so target savings scales proportionally with income. A higher permanent income therefore raises both realized wealth and the corresponding target, making their separation more difficult to interpret.

Relation to classic results on income risk and the certainty-equivalent benchmark. This paper takes a different perspective from Kimball (1990b) and Carroll and Kimball (1996). Their analysis studies how higher income risk increases the concavity of the consumption rule for a given type, raising the MPC relative to the certainty-equivalent benchmark at a fixed state. By construction, this approach abstracts from endogenous adjustments in target savings and from compositional changes in the cross-sectional distribution.[14] By contrast, my analysis compares effective allocations across types once savings adjust endogenously to risk. Stronger precautionary motives raise target savings, tilt consumption toward the future, and lower the MPC through greater intertemporal insurance. In this sense, certainty equivalence emerges here as the limit of high self-insurance rather than simply as a fixed-state benchmark. Appendix A uses compensating and equivalent precautionary premia to clarify this distinction.

[14] See also Commault (2022b) for a counterfactual approach based on compensating precautionary premia.

4 Stylized Facts

This section organizes the empirical evidence around the model's comparative statics restrictions. The focus is on income risk heterogeneity. The objective is not causal identification, but rather to assess whether the joint patterns emphasized by the mechanism are present in the data. The key restriction is that the incidence of saving motives across income levels—equivalently, the covariance between current or permanent income and income risk—governs the cross-sectional mapping between resources and MPCs. If risk is not correlated with, or is positively correlated with, income levels, the capacity to save and precautionary motives are expected to move together. Absent other unobserved constraint shocks, this alignment should preserve the standard aggregation logic: high-risk households are expected to be wealthier, and MPCs are expected to be monotonically decreasing both in risk and in wealth (or cash-on-hand).

If higher risk is paired with lower income, then the model's restrictions become essential, because high measured risk and low income may reflect preferences (e.g., impatience or risk tolerance). The SIM therefore delivers precise disciplining guidance.

If risk operates through precautionary motives and households are not effectively liquidity constrained, high-risk types should display higher expected consumption growth (less front-loading and higher target savings), and lower MPCs even at low wealth levels. If this is the case, the unconditional MPC–cash (or MPC–wealth) profile need not be monotone. By contrast, if high observed risk reflects preference-driven sorting (e.g., higher impatience), the model predicts lower target savings and therefore no increase in expected consumption growth; we would instead expect high MPCs among high-risk households. Thus, observing high risk together with low (or falling) consumption growth points to an explanation driven by preference heterogeneity rather than risk heterogeneity.

4.1 Empirical strategy

With this logic in mind, the empirical strategy proceeds in two steps. First, I measure household-level *uninsurable* income risk and study how its incidence varies across the distribution of resources. Second, I study how this measure affects consumption behavior.

Data. Both steps require panel data on income, consumption, and wealth, which are rarely jointly available in administrative sources. I therefore use the Bank of Italy's Survey of Household Income and Wealth (SHIW), a biennial rotating panel that jointly contains these variables. I focus on the Historical Sample 1998–2022. The baseline measure of labor income is defined as the sum of wage, self-employment income including real-estate income, and government transfers.[15] Consumption is defined as nondurable expenditures, while definitions of liquid and illiquid wealth components follow Kaplan

[15]Financial income is excluded, as it is realized less frequently (Kaplan et al., 2014). Real-estate income is included because more than 30% of Italian households regularly report rental income from buildings, especially at low levels of labor income, where renting often functions as a form of self-employment (Ministero dell'Economia e delle Finanze, Dipartimento delle Finanze, 2025). Omitting it would overstate household's income risk and mischaracterize households' ability to smooth consumption. The Online Appendix C assesses robustness to alternative income definitions.

et al. (2014). The sample is restricted to household heads aged 25–55 with at least three consecutive observations. The final sample consists of 9,779 individual-year observations. Details on the construction of the variables, sample selection, and descriptive statistics are provided in Appendix C.

Step 1. Income risk is measured following Arellano et al. (2022), using an individual-level coefficient of variation of *unpredictable* income:

$$CV_{i,t} \equiv \frac{\mathbb{E}\left(|Y_{i,t} - \mathbb{E}(Y_{i,t} \mid X_{i,t})| \mid X_{i,t}\right)}{\mathbb{E}(Y_{i,t} \mid X_{i,t})}. \tag{14}$$

Here $X_{i,t}$ proxies the household information set and includes rich micro and macro predictors (see Appendix C for details). The numerator measures the expected absolute deviation of income from its conditional expectation, given the information available at time t. Dispersion is therefore computed relative to predicted income, not around raw realized levels.

Two clarifications are important. First, this object differs from the standard coefficient of variation based on conditional standard deviations. It is a mean-absolute-deviation–based measure of relative dispersion, which is more robust to extreme values and does not rely on parametric variance assumptions. Second, although constructed from realized deviations, it is not a measure of raw income volatility. Because deviations are defined relative to a rich conditional expectation, predictable income dynamics are purged before dispersion is evaluated. The statistic therefore captures the magnitude of income innovations that are not predictable given $X_{i,t}$.

This distinction is central for the purposes of the analysis. The theoretical mechanism tested below hinges on forward-looking uncertainty that shapes target savings and intertemporal choices. Measures based on realized income changes or raw first differences would conflate predictable movements with genuine uncertainty. In contrast, $CV_{i,t}$ captures dispersion that remains after conditioning on observables, thereby providing a proxy for ex-ante income risk rather than realized instability.

Step 2. I use two different measures of MPCs out of transitory income shocks. First, I recover a robust version of the Blundell et al. (2008) estimator. Second, I present complementary results based on hypothetical self-reported MPCs available in selected waves of the SHIW. The BPP estimator, despite its well-known limitations discussed later, allows me to exploit the full sample, thereby increasing the effective sample size and extending the time coverage of the analysis.

The BPP implementation follows the usual two-step procedure. First, log income and log nondurable consumption are residualized with respect to year effects, demographic and geographic controls, and their interactions. Second, I construct first-difference residuals $\Delta y_{i,t}$ and $\Delta c_{i,t}$.

To control for potential sources of type heterogeneity beyond income risk (e.g. preferences), the BPP pass-through coefficient is estimated from the IV regression

$$\Delta c_{i,t} = \psi^{\varepsilon}_{BPP}\, \Delta y_{i,t} + u_{i,t},$$

where $\Delta y_{i,t}$ is instrumented with $\Delta y_{i,t+1}$ following the original BPP identification strategy. The regression is estimated using high-dimensional fixed effects that absorb household and year effects.[16]

Under the income growth process $\Delta y_{i,t} = \eta_{i,t} + \Delta\varepsilon_{i,t}$, the BPP estimator identifies the pass-through of transitory income shocks to consumption,

$$\psi^{\varepsilon}_{BPP} = \frac{\text{cov}(\Delta c_{i,t},\, \varepsilon_{i,t})}{\text{var}(\varepsilon_{i,t})} = \frac{\text{cov}(\Delta c_{i,t},\, \Delta y_{i,t+1})}{\text{cov}(\Delta y_{i,t},\, \Delta y_{i,t+1})}. \tag{15}$$

As shown by Commault (2022a), the average marginal propensity to consume out of transitory income shocks is given by the product between the BPP pass-through coefficient and the average propensity to consume,

$$\mathbb{E}[MPC] = \mathbb{E}\left[\frac{c_{i,t}}{y_{i,t}}\right] \psi^{\varepsilon}_{BPP}. \tag{16}$$

The original BPP estimator is often motivated under the assumption that log consumption follows a random walk. In my data, income growth $\Delta y_{i,t}$ displays a short-memory structure that is well approximated by a low-order moving-average process. Mechanically, this pattern is consistent with the standard permanent–transitory income decomposition. When income is written as the sum of a permanent component and a transitory shock, differencing income generates an MA(1) component in income growth. This occurs even if the transitory shock in income levels is serially uncorrelated. Empirically, the data support this short-memory structure. Higher-order autocovariances of consumption growth and cross-covariances with future income growth beyond one period are not statistically significant (Table C.2.1). This evidence indicates that income shocks dissipate quickly and that consumption growth does not display long-run persistence. Under these conditions, future income growth $\Delta y_{i,t+1}$ remains a valid instrument for $\Delta y_{i,t}$. Consistent with Kaplan and Violante (2010) and Commault (2022a), this allows for a robust implementation of the BPP estimator without imposing the random-walk assumption for consumption or assuming that households have full knowledge of their income process.

4.2 Four Facts on Income Risk Heterogeneity and Household Saving and Consumption Behavior

I now document four empirical facts about income-risk heterogeneity and its implications for saving and consumption behavior.

Fact 1. Income risk is heterogeneously persistent and negatively correlated with income. Figure 4 establishes two key properties of the estimated distribution of income risk. Panel (a) shows that the income risk is stable over time but highly dispersed.

[16] Because the identifying BPP moment condition, $E[\Delta y_{i,t+1}\, u_{i,t}] = 0$, is defined in first differences, absorbing fixed effects does not alter the orthogonality condition but only removes time-invariant heterogeneity across households. Nonetheless, Table C.2.4 shows that the results remain robust while dropping households fixed effects.

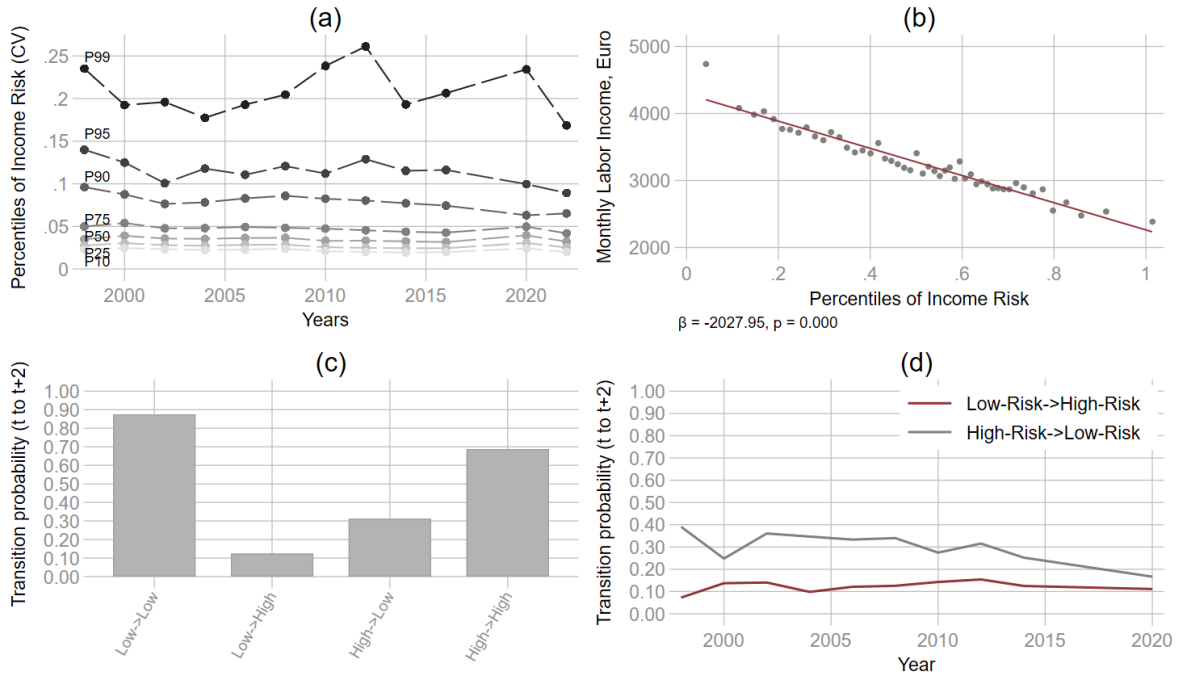


Figure 4: Income risk over time and its correlates

Notes: Income risk is measured as the coefficient of variation of household income and, when indicated, expressed as its within-year percentile rank. Panels reporting regressions control for age, sex, household size, education, marital status, region, and year fixed effects. The reported coefficient corresponds to income risk; standard errors are clustered at the household level, and statistical significance is reported in each panel. **Top panels.** The left panel shows the evolution of income-risk percentiles over time. The right panel reports binned scatter plots of monthly income against income-risk percentiles, conditional on controls. **Bottom panels.** Households are classified as *high-risk* or *low-risk* based on whether their income risk lies above or below the year-specific median. Transition probabilities are computed over two-year horizons using the biennial structure of the survey. The left panel reports average transition probabilities pooled across years; the right panel reports transition probabilities by initial year. Probabilities are relative frequencies and sum to one within each initial risk state.

Panels (c) and (d) show that households tend to remain above or below the median risk threshold over two-year horizons, indicating substantial persistence in relative risk positions. Panel (b) documents a strong negative relationship between income risk and labor income: moving from the 10th to the 90th percentile of income risk is associated with a decline in monthly income of roughly €2,000. This negative covariance between income levels and income risk is consistent with evidence for other countries (e.g., Mueller, 2017; Arellano et al., 2022; Broer et al., 2022) and implies that risk exposure is concentrated among lower-income households. Importantly, the persistence of relative risk ranks implies economically meaningful type heterogeneity that is not fully absorbed by fixed effects, motivating a model with persistent heterogeneity in precautionary motives.

Fact 2. Higher income risk is associated with lower liquid resources despite stronger saving motives. Figure 5 links income risk to saving and consumption behavior. Panel (a) shows that liquid assets relative to income decline monotonically with income risk. Panel (b) reports a similar sorting over the cash-on-hand distribution (defined as liquid asset plus monthly income). However, Panel (c) reports that there is no relationship between income risk and the probability of being constrained. The con-

straint indicator is constructed consistently with the standard incomplete-markets model, in which the MPC equals one when households hold zero liquid assets and consume their entire income. With ex-ante heterogeneity in income risk, the model predicts that a positive relationship with this probability would indicate low target savings. The absence of such a relationship therefore provides no evidence of low target savings.

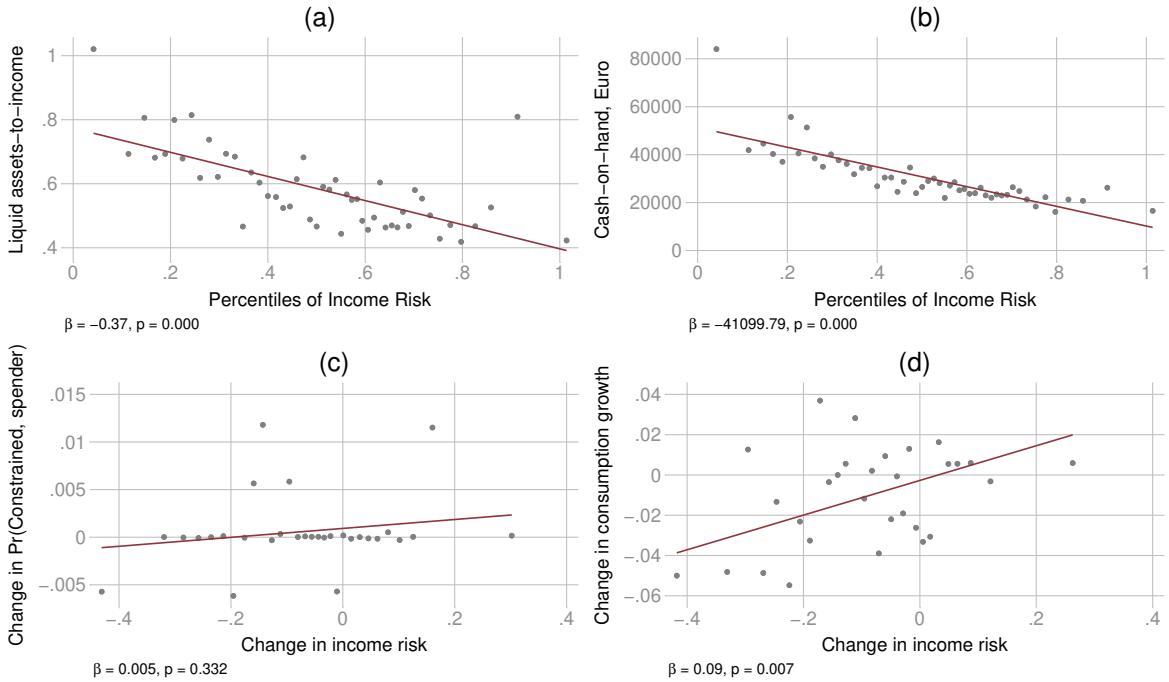


Figure 5: Income risk, savings and consumption growth

Notes: The figure reports the relationship between income risk and household outcomes. Panels (a) and (b) respectively show binned scatter plots of liquid assets-to-income ratios and cash-on-hand (defined as liquid asset plus monthly income) against the percentiles of income risk (measured as the coefficient of variation of income), controlling for demographics, region, and year fixed effects. Panel (c) reports the change in the probability of being constrained against the change in percentiles of income risk, controlling for demographics, region, and year fixed effects. A household is classified as constrained if (i) liquid financial assets are exactly zero and (ii) consumption equals income up to a rounding tolerance. Panels (d) report the relationship between the change income risk and the change in annualized two-year log consumption growth. The panel report results in first differences to remove time-invariant household heterogeneity.

Consistently, Panel (d) shows that higher income risk is associated with higher annualized consumption growth in first differences. This result is robust to fixed-effects estimates, indicating that within-household increases in income risk lead to higher subsequent consumption growth.[17] The results in Panels (c) and (d) suggest that low asset holdings among high-risk households are not primarily driven by lower patience or in general by preferences heterogeneity, but by lower income and therefore limited capacity to accumulate buffers. In short, consistent with the model's prediction, higher risk raises precautionary motives, but lower income reduces the ability to accumulate savings.[18]

[17]Specifically, I run a fixed effect regression of the form $gc_{i,t+2} = \beta\, risk_{i,t} + \gamma_t + X'_{i,t}\delta + \varepsilon_{i,t}$, where $gc_{i,t+2}$ denotes annualized consumption growth between t and $t+2$, $risk_{i,t}$ is the percentile rank of income risk, γ_t are year fixed effects, and $X_{i,t}$ includes demographic controls. The estimated coefficient on income risk is 0.021 ($P > |t| = 2.\%$, s.e. 0.009) without household fixed effects and 0.045 ($P > |t| = 2.1\%$, s.e. 0.019) when household fixed effects are included. Standard errors are clustered at the household level.

[18]These results also provide a plausible explanation for the classic literature that often finds no positive

Fact 3. High-risk households exhibit lower MPCs out of transitory shocks. The left panel of Figure 6 reports BPP transmission coefficients of transitory income shocks separately by income-risk type. High-risk households display significantly lower pass-through and implied MPCs than low-risk households, despite having lower income and lower liquidity. This result is consistent with the intertemporal self-insurance channel: households facing higher income risk have higher target savings, smoothing consumption more aggressively, and regardless of their cash-on-hand they exhibit lower MPCs.

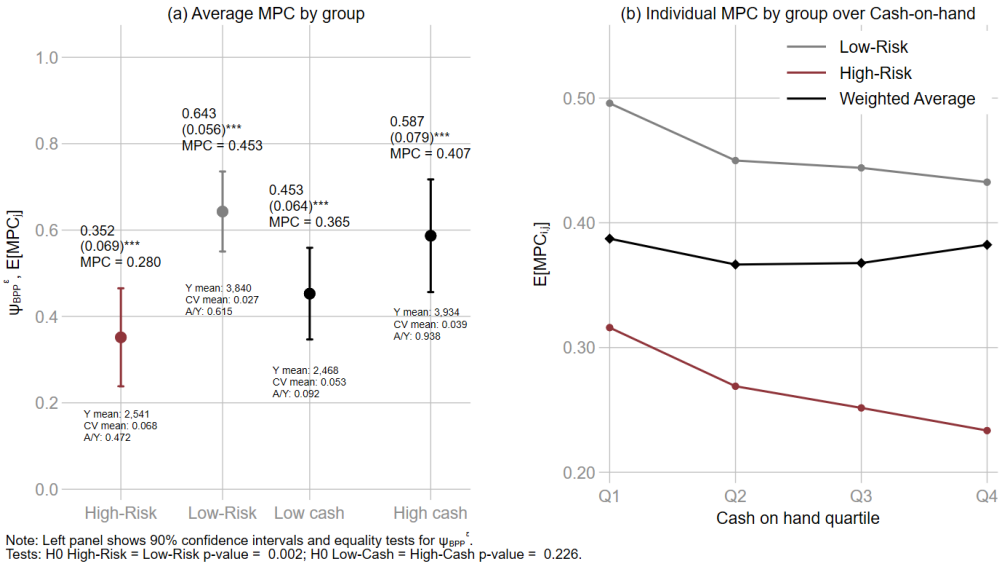


Figure 6: Heterogeneity in MPCs by Income Risk and Cash-on-Hand

Notes: The left panel reports the BPP transmission coefficient of transitory income shocks to consumption and the implied MPCs for household groups defined by income risk and cash-on-hand. *Low-risk* (*high-risk*) households have a CV below (above) the sample median; *Low cash* (*High cash*) households have cash-on-hand below (above) the median. The figure reports point estimates with 90% confidence intervals and p-values from tests of equality of the BPP transmission coefficients across groups. Coefficients are estimated via an IV specification with group-specific first stages and time and individual fixed effects. Instrument relevance is strong for both set of estimates: the Kleibergen–Paap Wald F statistics are 108.6 for the risk estimates and 324.3 for the cash estimates. Both models are exactly identified. The right panel shows implied within-group MPC profiles as a function of cash-on-hand, constructed by scaling the group-specific BPP transmission coefficients by individual APCs. Households are sorted into quartiles of cash-on-hand. Gray and red dots report quartile-averaged MPCs for low-risk and high-risk households, with solid lines denoting local linear polynomial smooths. The black line reports the mass-weighted average MPC within each cash-on-hand quartile, combining group-specific MPCs using within-quartile population shares.

Fact 4. The cross-sectional MPC–cash-on-hand relationship is mildly increasing, reflecting type-sorting. The right panel of Figure 6 presents a decomposition of MPCs. For each household i belonging to the risk group g, I construct an implied MPC distribution by scaling the group specific BPP estimate with the respective APCs distribution. That is: $\widehat{\text{MPC}}i = \widehat{\psi}^{\varepsilon}\text{BPP}(g) \cdot \text{APC}i$. Subsequently, I correlate these distributions with cash-on-hand quartiles. Within each risk group, the implied MPC decreases with cash-on-hand because APCi decreases with cash under the maintained mapping from

association between precautionary motives and savings (Pratt and Zeckhauser, 1987; Guiso et al., 1992; Bertaut and Haliassos, 1997; Dynan, 1993).

pass-through to MPC. Low-risk households' profiles are shifted upward, consistent with their higher estimated pass-through.

The mass-weighted average MPC across cash-on-hand quartiles is essentially flat. This pattern is consistent with the absence of statistically significant differences between the pooled low-cash and high-cash pass-through estimates reported in the left panel of Figure 6 and highlights the role of sorting. Although the APC declines with cash-on-hand for both groups, and level differences are accounted for by their respective average pass-through coefficients, the flat MPC profile reflects composition effects: higher cash quartiles contain a larger share of low-risk households, which display higher pass-through coefficients. For example, in the third quartile 60% of households are low-risk versus 40% high-risk, while in the fourth quartile 74% are low-risk compared with 26% high-risk. Of course, this evidence is suggestive and does not necessarily imply that pass-through coefficients are identical along the cash distribution. Clearly, allowing $\psi^{\varepsilon}_{\text{BPP}}$ to vary flexibly across both risk and cash-on-hand cells would sharpen the decomposition, but the available sample size does not permit precise estimation within risk-by-cash bins. Coarser partitions would mechanically pool variation along both dimensions, hindering the separation of true heterogeneity from sampling noise. Acknowledging these limitations, the exercise still provides a valid implementation of Proposition 5, suggesting that the aggregate MPC profile depends both on within-type MPCs and on how type shares vary with cash-on-hand.

Taken together, Figures 3–6 document persistent heterogeneity in income risk, a negative covariance between risk and resources, stronger precautionary motive and lower MPCs among high-risk households, and a composition-driven non-monotonic aggregate MPC–cash gradient. Collectively, these findings show that cash-on-hand alone cannot discipline the cross-section of MPCs, which are instead shaped by heterogeneity in target savings and, therefore, in intertemporal insurance. Before discussing the limitations of this findings and the necessary robustness checks, it is useful to present complementary evidence by exploiting the panel of self-reported MPCs available in the SHIW.

4.3 Evidence from Self-Reported MPCs

The last two waves of the SHIW (2020 and 2022) include a special module eliciting hypothetical consumption responses to transitory income shocks of a well-defined size and over a clearly specified spending horizon. Importantly, the question is identical in both waves:

> *"Suppose you unexpectedly received a sum of money equal to your households' monthly income. What percentage of this amount would you save and what percentage would you spend over the next 12 months?"*

I merge the 2020–2022 self-reported MPC module with the main historical sample (1998–2022). The merged file contains 957 household–wave observations (360 in 2020 and 597 in 2022) for which both self-reported MPCs and BPP-implied MPC measures

are available.[19] As shown in Figure C.3.2, this restricted sample displays properties remarkably similar to the baseline sample, which is important for the comparability of responses and estimates. The qualitative patterns remain unchanged: higher income and cash-on-hand are associated with lower income risk, but the gradients are steeper in 2020–2022. This suggests that marginal increases in risk are accompanied by larger reductions in available resources in the more recent period, possibly reflecting the heightened uncertainty during the COVID episode. Figure C.3.1 in the appendix reports the distribution of self-reported MPCs, which appears remarkably stable across the 2020 and 2022 waves. This stability is reassuring given the pandemic context and speaks to the internal consistency of the survey instrument. Table 1 estimates a single baseline specification,

$$MPC_{it} = \beta^{\mathcal{CH}}\mathcal{CH}_{it} + \beta^{\mathcal{CV}}\mathcal{CV}_{it} + X_{it}'\Gamma + \alpha_i + \lambda_t + \varepsilon_{it}, \tag{17}$$

where $\mathcal{CH}_{it}$ and $\mathcal{CV}_{it}$ denote the within-year percentile of cash-on-hand and income risk, respectively. The vector X_{it} collects controls, α_i are household fixed effects, and λ_t are year fixed effects. Identification comes from within-household variation over the 2020–2022 panel. Columns differ in whether specifications absorb only household fixed effects or additionally absorb fixed effects for risk bins or cash bins.

The estimates in columns (1) and (2) illustrate that cash and risk effects can be confounded when considered in isolation even in a panel setting. Column (1), which relates MPCs to cash-on-hand without conditioning on risk heterogeneity, yields a positive and statistically insignificant average cash slope. Column (2) instead relates MPCs to income risk and finds a consistent negative association, but statistically insignificant.

Column (3) instead, estimates the within-household cash effect while absorbing risk-bin fixed effects, thereby removing level differences across risk bins. In this specification the cash slope becomes negative and statistically significant, consistent with concavity of the consumption rule within risk types. Symmetrically, column (4) estimates the within-household risk effect while absorbing cash-bin fixed effects and yields a negative and statistically significant coefficient.[20] Column (5) allows the MPC–cash slope to differ by risk type. For high-risk households (the reference group), the estimated cash slope is negative (-0.355). For low-risk households the interaction coefficient of 0.878 implies an overall positive slope. Hence the MPC–cash gradient differs across risk types. Accordingly, the predicted MPCs evaluated at the average (type-representative) cash percentiles of low-risk (59th) and high-risk (42nd) households are 0.61 and 0.35, respectively.

Taking stock. The self-reported MPCs confirm the BPP evidence. Overall, the empirical exploration shows that (i) risk is persistent and negatively related to household

[19] Similar questions appear in the 2010 and 2016 SHIW waves used by Jappelli and Pistaferri (2020). However, the spending horizon is not explicitly defined, complicating the interpretation of responses and limiting comparability with the later questions. Comparability with the full sample is further constrained by the six-year gap between waves and the rotating-panel structure of the survey, which yields a small overlapping sample (about 60 households). In that sample, income is positively correlated with risk and MPCs decline sharply with cash-on-hand, consistent with Jappelli and Pistaferri (2020).

[20] Because the panel contains only two waves, the combination of household and bin fixed effects absorbs a large share of variation and reduces the usable identifying variation, which is reflected in the smaller estimation samples in columns (3) and (4).

resources; (ii) higher risk is associated with higher expected consumption growth, consistent with stronger precautionary motives and higher target savings; (iii) consistent with the theoretical mechanism, high-risk households display lower MPCs despite lower resources; and (iv) the aggregate MPC–cash-on-hand profile can be flat or even increasing because high-MPC (low-risk) households are disproportionately represented at higher cash levels.

Table 1: Cash and Curvature Effects

	(1) MPC-cash	(2) MPC-risk	(3) Cash Effect	(4) Risk Effect	(5) MPC-cash shift
Percentiles of cash	0.193		-0.631***		-0.355
	(0.187)		(0.139)		(0.288)
Percentiles of risk		-0.358		-0.818***	
		(0.214)		(0.243)	
Percentiles of cash × low-risk					0.878**
					(0.371)
Predicted MPC at 59th percentile of cash					0.607***
					(0.044)
Predicted MPC at 42nd percentile of cash					0.354***
					(0.052)
Observations	200	200	106	136	200
Within R^2	0.15	0.23	0.95	0.94	0.25
Controls	Yes	Yes	Yes	Yes	Yes
Individual fixed effects	Yes	Yes	Yes	Yes	Yes
Risk fixed effects			Yes		
Cash fixed effects				Yes	

Notes: Column (1) estimates the within-household association between MPC and cash rank, abstracting from risk heterogeneity. Column (2) estimates the within-household association between MPC and income risk, abstracting from cash heterogeneity. Column (3) estimates the within-household cash effect while absorbing risk-bin fixed effects. Column (4) estimates the within-household risk effect while absorbing cash-bin fixed effects. Column (5) allows the MPC–cash slope to differ by risk type. Reported below are the predicted MPCs evaluated at the average (type-representative) cash percentiles of low-risk (59th) and high-risk (42nd) households, respectively. Their difference is statistically significant ($P < 0.01$). Standard errors are clustered at the household level and reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Low-risk is an indicator for being below the median of the CV distribution. MPC–cash regressions control for demographics, labor-market characteristics, and saving attitudes. MPC–risk regressions additionally control for income expectations, access to informal credit in case of emergency, south-regions×age interactions, and a 2020 dummy.

Discussion. The empirical evidence presented so far is best interpreted as disciplining rather than identifying the mechanism. The goal was not to establish a causal effect of income risk on MPCs, but to assess whether the joint restrictions implied by the model are visible in the data. Taken together, these stylized facts are difficult to reconcile with a single-type SIM in which cash-on-hand is a sufficient statistic for MPC heterogeneity, but arise naturally under standard preferences and rational expectations once ex-ante heterogeneity is modeled in a way that identifies the intertemporal margin of self-insurance.

These empirical results should nevertheless be interpreted within clear boundaries. A first concern relates to the measurement of risk. The CV is constructed to purge predictable movements in earnings by conditioning on a rich information set. However, as with any residual-based statistic, it may still reflect short-lived fluctuations, rare events,

or measurement error. In particular, one may worry that it captures noisy realized volatility rather than persistent ex-ante dispersion. Two pieces of evidence mitigate this concern. First, exploiting the panel dimension (Figure 4), the risk classification is highly persistent over time, indicating that the measure reflects stable heterogeneity in exposure to income risk rather than temporary instability. Households do not randomly switch between low- and high-risk states across waves. Second, when standard ex-post volatility measures are computed, households above the median of the CV distribution exhibit systematically higher variability in both the permanent and transitory components of income (Appendix D.3). This also addresses a related concern about the normalization embedded in the CV statistic: because the CV divides dispersion by predicted income, identical absolute innovation dispersion could mechanically generate higher measured risk at lower income levels. However, the decomposition in Appendix D.3 is based on absolute variances and does not normalize by income, indicating that the ranking induced by the CV statistic reflects genuine differences in income risk rather than a mechanical consequence of the normalization. This interpretation is further supported by the fact that households with temporary employment contracts are disproportionately located in the upper percentiles of the CV distribution. Of course, the measure does not separately identify permanent and transitory components, nor does it capture tail risk or higher-order moments that may matter for precautionary behavior (e.g., Ghosh and Theloudis, 2024). Addressing these dimensions would require a structural specification beyond the scope of the present analysis.

A second concern is that the consumption-growth gradients reflect predictable income dynamics rather than precautionary behavior. In a biennial setting, high-risk households could simply be those recently hit by shocks, so that subsequent consumption growth captures mechanical income recovery or labor market transitions. The CV addresses this possibility directly, as it is conditional on a rich information set—including past income realizations, employment history, contract type, and macro conditions—thereby purging predictable income movements before risk is constructed. If predictable rebounds were driving the gradient, conditioning on observables would substantially attenuate it. Instead, Figure 5 shows that the positive association between risk and subsequent consumption growth persists both in cross-sectional specifications and in panel regressions with household fixed effects, where increases in measured risk within the same household are followed by higher consumption growth, conditional on observables. This within-household evidence is difficult to reconcile with a purely mechanical recovery story.

A related concern is that measured income risk may proxy for preference heterogeneity rather than exposure to risk. For example, households selecting into high-risk jobs may be less risk-averse or more impatient. But this interpretation has a clear empirical implication: if higher measured risk reflected weaker precautionary motives, it should be associated with lower, not higher, subsequent consumption growth. The data reject this prediction, and are consistent with the view that the CV captures meaningful differences in exposure to income risk. This interpretation is further strengthened by the joint behavior of expected consumption growth and MPCs. Consider first a purely within-type interpretation, in which low wealth reflects only movement along a common consumption rule. Under this logic, households with higher expected consumption growth should lie

on a portion of the policy function associated with higher, not lower, MPCs. The data show the opposite. Consider next an across-type interpretation in which low resources are driven by preferences rather than by a limited ability to save. If high-risk, low-resource households were characterized by greater impatience or lower risk aversion, this explanation would again predict higher, not lower, MPCs. Importantly, Figure 5 also shows that higher risk is not associated with a greater probability of operating "at" the constraint. Taken together, these results indicate that low asset holdings among high-risk households are not primarily driven by weaker precautionary motives or greater impatience. Rather, higher risk raises target savings, while lower income limits the ability to accumulate large wealth positions. Consistent with the model, intertemporal insurance improves even when realized wealth remains low; accordingly, these households display lower MPCs.

A third concern is of course about the interpretation of the pass-through estimates. Table C.2.4 in the Online Appendix shows that the results reported in Figure 6 are robust to a broad set of alternative specifications, including different definitions of risk, income, resources, and consumption, as well as to alternative sample periods and age restrictions. However, beyond these sample discretionary choices, the main concern is that with measurement errors, and if groups differ in shock persistence, signal-to-noise ratios, or advance information, the BPP estimator may mechanically generate different coefficients even when behavior is identical, see Pedroni et al. (2022). To address this, in Section 5 I apply the BPP estimator to model-simulated data at SHIW frequency. The implied MPC moment constructed from the BPP pass-through coefficient (i.e. the product of the pass-through and the APC) remains close to the model-implied MPC, defined as the local derivative of consumption with respect to cash-on-hand holding expectations fixed, and thus supporting the identification of the mechanism. To further take into account the possible contamination of advance information in the empirical estimates, table C.2.2 reports estimates under alternative assumptions regarding agents' information sets and the magnitude of measurement error, explicitly incorporating evidence from the SHIW documented by Kaufmann and Pistaferri (2009). Using income expectation measures available in the 1995 and 1998 SHIW waves, they show that transitory income variation can be decomposed into an anticipated component (approximately 50%), an unanticipated component (20%), and measurement error (30%). The main qualitative results are preserved: pass-through for low-risk households remains higher and statistically significant, and the conclusions are robust to more or less stringent assumptions about anticipation and measurement error. Additional diagnostic tests clarify this result: Table C.2.1 shows that, in the main sample, the covariance between consumption growth and future income growth becomes negligible beyond one period ahead, indicating a limited anticipation horizon and mitigating concerns about correlations between advance information and estimated income risk. In this context, these concerns are further attenuated by the strong explanatory power of the first stage in the CV estimates (Table D.1.1), where the information set captures 98% of the variation. In fact, when computing the standard deviations of permanent and transitory shocks separately by risk group (Figure D.3.1), clear differences emerge, particularly in the dispersion of permanent shocks: the standard deviation is 0.27 for high-risk households and 0.21 for low-risk households. By contrast, the gap

in transitory shock volatility is more limited (0.41 for high-risk versus 0.38 for low-risk). These patterns confirm that the coefficient of variation effectively captures overall income risk, with especially strong exposure to its permanent component. Finally, Table C.2.3 in the Online Appendix replaces the permanent component with a persistent AR(1) process. The results remain robust across a wide range of autoregressive parameterizations. Although these checks cannot fully eliminate all potential sources of bias—particularly those arising from advance information—the model-simulated implementation of the BPP estimates, together with the aligned evidence from the hypothetical MPC module (which explicitly specifies the spending horizon and shock size), provides meaningful validation of the mechanism proposed in this paper. In particular, the ranking of MPCs across risk groups remains stable across specifications.

A fourth concern is that the sorting over the cash quartiles of low-risk households here may be sample- and period-specific. The strength of the income–risk covariance, the prevalence of liquidity constraints, and the mapping between risk exposure and insurance opportunities can vary across institutional settings, time periods, and measurement choices. In environments where high risk is concentrated among households that are frequently at the borrowing constraint (e.g., due to unemployment spells or limited access to informal credit), liquidity effects may dominate and the MPC–risk relationship may become non-monotone, consistent with the robustness discussion in Appendix B. Hence, the results do not imply that affluent households *always* have higher MPCs; rather, they show that this pattern can arise once precautionary motives are separated from liquidity effects. More generally, the paper shows that cash-on-hand is not necessarily a sufficient statistic for MPC heterogeneity, and that reduced-form MPC–wealth gradients are inherently compositional objects.

5 Quantitative Evidence

This section assesses whether a SIM framework augmented with the joint heterogeneity structure emphasized in the data can replicate the key cross-sectional empirical patterns documented in Figure 6 and therefore provides a practical validation for Propositions 4 and 5.

Environment. The economy is populated by a unit measure of infinitely lived households indexed by $i \in [0,1]$. Time is discrete. Guided by the evidence of persistent heterogeneity in income risk and its negative covariance with income levels (Figure 3), ex-ante heterogeneity is introduced by partitioning the population into two permanent types $j \in \{L, H\}$: a share $\pi = 1/2$ of high-risk households (H) and a share $1-\pi$ of low-risk households (L). Within each type, households are heterogeneous in asset holdings a and idiosyncratic labor income y.

A household of type j solves

$$V_j(a_{t-1}, s_t) = \max_{c_t, a_t} \; u(c_t) + \beta \, \mathbb{E}[V_j(a_t, s_{t+1}) \mid s_t] \tag{18}$$

$$\text{s.t.} \qquad a_t + c_t = (1+r)a_{t-1} + \psi_{y,j} \, y_t(s_t), \tag{19}$$

$$a_t \geq 0, \tag{20}$$

where r is the exogenous interest rate, $\beta \in (0,1)$ is the discount factor, and $\psi_{y,j}$ is a type-specific permanent-income factor (skill premium), capturing the capacity to save. Consistent with the evidence that low-risk households have higher income levels, I impose $\psi_{y,L} > \psi_{y,H}$.

Preferences are CRRA:

$$u(c) = \begin{cases} \dfrac{c^{1-\gamma}}{1-\gamma}, & \gamma \neq 1, \\ \log(c), & \gamma = 1, \end{cases} \tag{21}$$

with relative risk aversion $\gamma > 0$.

Income process. Idiosyncratic income is the exponential of a persistent and a transitory component:

$$\log y_t = z_t + \varepsilon_t, \tag{22}$$

$$z_t = \rho z_{t-1} + \eta_t, \tag{23}$$

where $\varepsilon_t \sim \mathcal{N}(0, \sigma^2_{\varepsilon,j})$ and $\eta_t \sim \mathcal{N}(0, \sigma^2_{\eta,j})$ are i.i.d., mutually independent shocks. Risk heterogeneity is captured by allowing $(\sigma_{\varepsilon,j}, \sigma_{\eta,j})$ to differ across types, with high-risk households facing larger variances.[21]

MPC. The MPC is computed as the local derivative of the stationary consumption policy with respect to individual cash-on-hand x:

$$\mathrm{MPC}_j(x,s) \equiv \frac{\partial c_j(x,s)}{\partial x}. \tag{24}$$

Here $c_j(x,s)$ denotes the stationary consumption policy for a household of type j as a function of cash-on-hand x and the income state s. The MPC is therefore the partial derivative with respect to x, holding the income state s and household type j fixed, and thus holding expectations about future income constant.[22]

Parametrization. Consistent with the structure of the SHIW, the model is calibrated at a biennial frequency. To discipline income dynamics, I estimate an *auxiliary* covariance-structure model for residual log earnings separately for low-risk and high-risk households using a minimum-distance procedure. The auxiliary specification is a random-walk permanent component with an AR(1) transitory component and is used only to match the

[21] Note that income risk is negatively related to the permanent skill premium ψ. The discretized income states are normalized such that the stationary mean of income equals one, $E[y] = 1$. This normalization ensures that variations in the variance parameters affect only income risk and not the average level of current income.

[22] Other MPC concepts such as the MPC out of a discrete transfer or intertemporal MPCs are consistent with the evidence presented below and remain robust under this textbook definition. They are available upon request and not reported here, as their interpretation would require substantially more institutional and timing context than is necessary for the general point of this exercise.

empirical autocovariance structure of biennial income growth in the data. Details are presented in Appendix D.3. The parameters of the income process in (33), $(\sigma_{\eta,j}, \sigma_{\varepsilon,j})$, are chosen so that the model reproduces the same targeted second moments of income growth implied by the auxiliary estimates, specifically $Var(\Delta \log y_t)$ and $Cov(\Delta \log y_t, \Delta \log y_{t+1})$, for each type j. This choice is closely related to the identification strategy underlying the BPP pass-through estimates, which are themselves functions of the covariance structure of income growth. Matching these moments therefore ensures that the income process feeding the quantitative model is consistent with the empirical measures estimated in the data. The reported parameters are expressed as biennial standard deviations (Table 2). Consistent with the patterns highlighted in Figure 4, high-risk households exhibit greater dispersion in both the persistent and transitory components. Having disciplined income risk for the two groups, the model is used to match the estimates in the left panel of Figure 6. The discount factor β is restricted to be common across groups, ruling out preference heterogeneity as a source of MPC differences. Conditional on the estimated income processes, β and the relative income shifter $\psi_{y,H}$ are chosen to match the asset moments in Table 2.[23]

Table 2: Parameter Values: Biennial Calibration

Parameters	**Value**	**Description**	**Target**
Common parameters			
r	3.5%	Interest rate	–
$\theta = 1/\gamma$	0.40	Intertemporal Elasticity of Substitution	Empirical Evidence**
β	0.761	Implied Biennial Discount factor	$A_L/Y_L = 0.615, A_H/Y_H = 0.472$
Income process			
ρ_{zj}	0.9851	Persistence of AR(1) component	Aggregate Gini index: 0.778
σ_{η_L}	0.2115	SD Persistent shock, low-risk	SHIW estimates
σ_{η_H}	0.2715	SD Persistent income, high-risk	SHIW estimates
σ_{ε_L}	0.3845	SD of Transitory income, low-risk	SHIW estimates
σ_{ε_H}	0.4130	SD Transitory income, high-risk	SHIW estimates
Permanent Income			
$\psi_{y,L}$	1	Permanent skill premium, low-risk	Internally calibrated
$\psi_{y,H}$	0.402	Implied Permanent skill premium, high-risk	$A_H/Y_H = 0.472$

Notes: The model period is biennial. The interest rate is calibrated as a net real rate per model period, $r = 3.5\%$, so that $R = 1 + r = 1.035$ enters the Euler equation. All other parameters and targets are expressed in biennial units. The implied discount factor is 0.761 at the biennial frequency, corresponding to 0.873 annually and 0.967 quarterly. ** Using data from the European Central Bank, Sciacchetano (2024) estimates an intertemporal elasticity of substitution (IES) for Italy in the range of 0.4–0.5.

In addition to the MPC moments reported in Figure 6, the calibration imposes discipline on the cross-sectional distribution of liquid wealth and cash-on-hand. Specifically, the model matches the average liquid-asset-to-income ratios for high-risk and low-risk households (0.47 and 0.61, respectively), thereby aligning the two within-type wealth distributions with their empirical counterparts. The model also replicates the ranking of households with zero liquid assets, despite this moment not being explicitly targeted in

[23] Because the stationary mean of income equals one, $E[y] = 1$, the model wealth-to-income ratio equals average assets. The parameter $\psi_{y,H}$ captures differences in permanent income (accumulation capacity), while cross-group asset differences reflect precautionary motives induced by income risk.

the calibration. In the data, 6% of low-risk households hold zero liquid assets, compared with 3% in the model. For high-risk households, the corresponding shares are 11% in the data and 7% in the model. Consistent with these moments, the calibrated economy exactly matches overall wealth inequality: the model reproduces the SHIW Gini coefficient (0.778 in both the model and the data) and closely tracks the Lorenz curve.[24]

Given this wealth discipline, the MPC implications are discussed next. This wealth discipline is important because it constrains the extent to which MPC heterogeneity can be attributed to mechanical differences in resources.

Results. The top panels of Figure C.3.1 report the model analog of Figure 6. The left panel shows that the model reproduces the risk-gradient in MPCs. Low-risk households are wealthier on average by construction and display an average MPC of 0.40 in the model (0.45 in the data). Poorer and high-risk households display an average MPC of 0.31 in the model, close to 0.28 in the data. Thus, the model successfully delivers MPC differences across risk groups unrelated to cash-on-hand holdings, consistent with Proposition 4 and the mechanism based on heterogeneous saving motives.

Across liquidity groups, the model generates an MPC of 0.41 for households below the median cash-on-hand, broadly consistent with the empirical estimate (0.37), but it understates the MPC of high-cash households (0.30 in the model versus 0.41 in the data). This discrepancy reflects the structure of the simple one-asset environment, whereby cash-on-hand gives total resources. In the model, conditional on belonging to the top cash quartile, households are substantially wealthier than their empirical counterparts. This stronger mapping between liquidity and wealth mechanically amplifies intertemporal smoothing at the top and pushes down the predicted MPC. Nonetheless, the MPC remains relatively high in the model for two main reasons. First, the model is only partially calibrated to liquid assets, second, it successfully replicates the direction and monotonicity of type sorting across cash-on-hand quartiles.[25]

The top-right panel illustrates the qualitative pattern of the MPC across cash-on-hand quartiles. These moments are not directly targeted. In the data, the BPP is estimated at the group level (low- and high-risk), and heterogeneity across cash-on-hand quartiles arises from interacting these group-specific pass-through estimates with APC distributions that vary along the cash distribution. Hence, differences across quartiles are driven by heterogeneity in APC rather than by variation in the pass-through itself.

[24]Note that the parameter ρ_j is assumed to be common across types and is calibrated to match wealth inequality. The implied value corresponds to a persistent AR(1) income component rather than a strictly permanent one. This modeling choice is not central for the results: empirically, MPC estimates remain quantitatively similar when the BPP estimator is implemented under AR(1) income dynamics. See Table C.2.3 in the Online Appendix.

[25]These shares are intentionally left untargeted to provide an external validation of the mechanism. In the data, the high-risk share declines from 0.60 in the lowest cash quartile to 0.25 in the highest, whereas in the model it declines from 0.76 to 0.31. Conversely, for the low-risk group, the share increases from 0.40 to 0.75 in the data and from 0.24 to 0.69 in the model. In the model, within-quartile MPCs for low-risk households always exceed those for high-risk households. The model-implied values (low-risk vs. high-risk) are: Q1: 0.58 vs. 0.41; Q2: 0.46 vs. 0.28; Q3: 0.38 vs. 0.25; and Q4: 0.30 vs. 0.21. These are consistent with values in the data: Q1: 0.496 vs. 0.316; Q2: 0.450 vs. 0.269; Q3: 0.444 vs. 0.252; and Q4: 0.433 vs. 0.233.

The sorting mechanism, however, explains the aggregate patterns. Consistent with this channel, the model reproduces the broadly flat gradient observed over the first three quartiles—where model and data are most closely aligned in terms of liquidity—while declining only to 0.29 in the fourth quartile.[26]

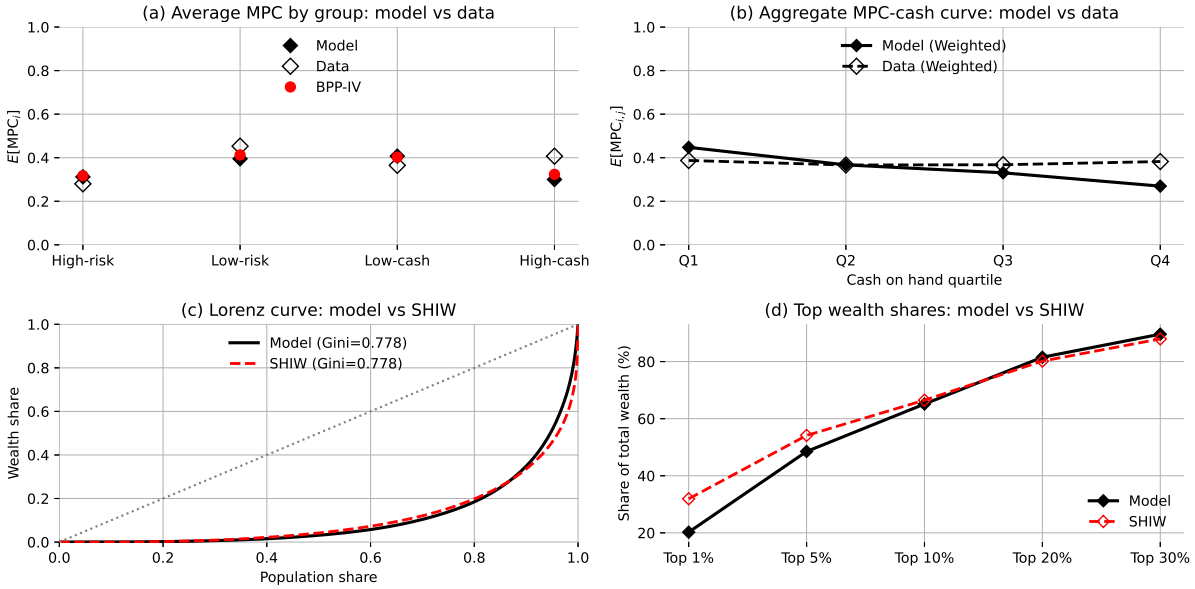


Figure 7: MPC cross-sectional heterogeneity and inequality: model vs data.

Notes: Panel (a) contrasts average MPCs across groups (high-risk, low-risk, low-cash, high-cash) in the model (equation 24), in the data, and from BPP estimates out of simulated data from the model. Panel (b) reports weighted aggregate MPCs by cash-on-hand quartile (Q1–Q4), comparing model-implied values with weighted empirical moments reported in Figure 6. Panel (c) displays the Lorenz curves for liquid wealth in the model and in the SHIW data, together with the associated Gini coefficients. Panel (d) reports top liquid wealth shares (top 1%, 5%, 10%, 20%, and 30%) in the model and in the SHIW data.

Alignment of theoretical and empirical MPC definitions. To assess whether the empirical patterns reflect the proposed structural behavior rather than estimator properties, I apply the BPP estimator to model-simulated data at SHIW frequency using the calibrated economy. The implied IV-MPCs closely track the model's structural MPCs. For low-risk households, the structural MPC of 0.40 maps into an IV-MPC of 0.41; for high-risk households, 0.31 maps into 0.32. Similarly, for low-cash households the structural MPC is 0.41 and IV-MPC is 0.40, while for high-cash households the IV-MPC (0.32) remains close to the structural value (0.30). The close correspondence indicates that the empirical MPC patterns are not driven by the estimator properties and are consistent with the model's mechanism.

Taken together, the calibrated economy (i) preserves realistic wealth dispersion, (ii) reproduces untargeted cross-sectional sorting across liquidity bins, and (iii) broadly matches the degree of MPC heterogeneity observed in the data. Combined with the empirical evidence, these results indicate that the mechanism proposed in this paper arises in empirically plausible regions of the parameter space.

[26]The aggregate asset-to-income ratios by cash-on-hand quartile are 0.04, 0.16, 0.48, and 14.15 in the model, compared with 0.02, 0.19, 0.44, and 1.56 in the data.

6 Conclusions

This paper shows that MPC heterogeneity reflects two self-insurance margins: a liquidity margin that matters near borrowing constraints and an intertemporal margin that operates through target savings and forward-looking plans in the interior. When saving motives and accumulation capacity co-move, standard one-type incomplete-markets models confound these margins, so resources remain close to sufficient statistics for MPCs and precautionary motives are hard to identify separately from liquidity.

The paper isolates the condition that breaks this alignment: a negative incidence between target-saving motives and the ability to accumulate wealth. Under this misalignment, households can be poor yet strongly insured intertemporally, while affluent households can have weak saving motives and high MPCs despite being unconstrained. Hence, under standard preferences and rational expectations, cross-sectional MPC–resource profiles can be non-monotone even if within-type MPCs decline with resources, due to sorting across types.

Consistent with this mechanism, Italian micro data show persistent income-risk heterogeneity that is negatively related to income, higher expected consumption growth for riskier households (suggesting higher target savings), and lower MPCs among high-risk households despite lower resources. A parsimonious quantitative model disciplined to match these joint patterns reproduces the observed MPC–cash profile while matching liquid-wealth inequality. The broader implication is that identifying MPC heterogeneity requires information on both resources and determinants of target savings: liquidity and precautionary motives need not coincide in the cross-section.

A natural next step is to embed this joint heterogeneity in general-equilibrium HANK models and study its implications for fiscal and monetary transmission. This requires new micro-to-macro objects capturing business-cycle covariances between saving ability, target wealth, and MPCs, beyond the standard covariances tied to constrained households (e.g., Bilbiie (2018); Patterson (2023)). The evidence here already suggests that policy transmission need not necessarily operate through borrowing constraint channels. When precautionary motives and resources are misaligned, self-insurance and composition among affluent households can be first-order for aggregate dynamics. Consistent with this view, Challe and Savoia (2026) show sizable implications for monetary transmission in a US-calibrated HANK model with search and matching and an endogenous negative relationship between permanent income and separation risk.

More broadly, the paper provides a structural account of how unobserved heterogeneity operates within standard models and rationalizes MPC patterns often attributed to behavioral departures. This is important because policymakers ultimately require frameworks grounded in observable frictions. The contribution is not to argue that a particular model best captures household heterogeneity—any framework has limitations and direct comparisons are often misleading—but rather to offer a novel perspective on its sources. In particular, the paper points to the role of affluent households, sorting mechanisms, and constraints on savings accumulation, as well as their interaction with target savings. Advancing research along these dimensions is key to better characterizing household heterogeneity, plausibly quantifying its aggregate implications, and understanding its role in

shaping aggregate demand dynamics and the transmission of macroeconomic policy.

References

Aguiar, M., Bils, M., and Boar, C. (2024). Who are the hand-to-mouth? *The Review of Economic Studies*.

Aguiar, M. A., Bils, M., and Boar, C. (2020). Who are the hand-to-mouth? *National Bureau of Economic Research, Inc*, (26643).

Aiyagari, S. (1994). Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics*, 109(3):659–684.

Amberg, N., Jansson, T., Klein, M., and Picco, A. R. (2022). Five facts about the distributional income effects of monetary policy shocks. *American Economic Review: Insights*, 4(3):289–304.

Andersen, A. L., Johannesen, N., Jørgensen, M., and Peydró, J.-L. (2023). Monetary policy and inequality. *The Journal of Finance*, 78(5):2945–2989.

Andre, P., Flynn, J. P., Nikolakoudis, G., and Sastry, K. (2025). Quick-fixing: Near-rationality in consumption and savings behavior. NBER Working Paper 33464, National Bureau of Economic Research. Revised September 2025.

Arellano, M., Bonhomme, S., De Vera, M., Hospido, L., and Wei, S. (2022). Income risk inequality: Evidence from spanish administrative records. *Quantitative Economics*, 13(4):1747–1801.

Auclert, A., Bardóczy, B., Rognlie, M., and Straub, L. (2021). Using the sequence-space jacobian to solve and estimate heterogeneous-agent models. *Econometrica*, 89(5):2375–2408.

Auclert, A., Rognlie, M., and Straub, L. (2020). Micro jumps, macro humps: Monetary policy and business cycles in an estimated hank model. Manuscript.

Auclert, A., Rognlie, M., and Straub, L. (2025). Fiscal and monetary policy with heterogeneous agents. *Annual Review of Economics*.

Berger-Thomson, L., Chung, E., and McKibbin, R. (2009). Estimating marginal propensities to consume in australia using micro data. *Reserve Bank Of Australia*.

Bergholt, D., Furlanetto, F., and Mori, L. (2025). Consumption inequality, household risks, and the business cycle. Technical Report 18/2025, Norges Bank.

Bertaut, C. C. and Haliassos, M. (1997). Precautionary portfolio behavior from a life-cycle perspective. *Journal of Economic Dynamics and Control*, 21(8-9):1511–1542.

Bilbiie, F. O. (2018). Monetary policy and heterogeneity: An analytical framework. *Available at SSRN 3106805*.

Bilbiie, F. O., Galaasen, S., Gürkaynak, R. S., Maehlum, M., and Molnar, K. (2024). Hanksson. Technical Report 20090, CEPR.

Blundell, R., Pistaferri, L., and Preston, I. (2008). Consumption inequality and partial insurance. *American Economic Review*, 98(5):1887–1921.

Boehm, J., Fize, E., and Jaravel, X. (2025). Five facts about mpcs: Evidence from a randomized experiment. *American Economic Review*, 115(1):1–42.

Boutros, M. (2021). Bounded intertemporal rationality and the marginal propensity to consume. *Duke University*.

Broer, T., Kramer, J., and Mitman, K. (2022). The curious incidence of monetary policy across the income distribution. *Sveriges Riksbank Working Paper*, (416).

Bunn, P., Le Roux, J., Reinold, K., and Surico, P. (2018). The consumption response to positive and negative income shocks. *Journal of Monetary Economics*, 96(C):1–15.

Caplin, A., Gregory, V., Lee, E., Leth-Petersen, S., and Sæverud, J. (2023). Subjective earnings risk. NBER Working Paper w31019, National Bureau of Economic Research.

Carroll, C. (2001). A theory of the consumption function, with and without liquidity constraints. *Journal of Economic Perspectives*, 15(3):23–45.

Carroll, C., Holm, M., and Kimball, M. (2021). Liquidity constraints and precautionary saving. *Journal of Economic Theory*, 195:105–276.

Carroll, C. and Kimball, M. (1996). On the concavity of the consumption function. *Econometrica*, 64(4):981–992.

Carroll, C., Slacalek, J., Tokuoka, K., and White, M. (2017). The distribution of wealth and the marginal propensity to consume. *Quantitative Economics*, 8(3):977–1020.

Carroll, C. D. (1997). Buffer-stock saving and the life cycle/permanent income hypothesis. *Quarterly Journal of Economics*, 112(1):1–56.

Challe, E. and Savoia, E. (2026). Monetary transmission with heterogeneous unemployment risk. *Mimeo*.

Colarieti, R., Mei, P., and Stantcheva, S. (2024). The how and why of household reactions to income shocks. Working Paper w32191, National Bureau of Economic Research.

Colarieti, R. and Monacelli, T. (2023). H-hank: A heteroskedastic heterogeneous-agent new keynesian model. *Unpublished Manuscript*.

Commault, J. (2022a). Does consumption respond to transitory shocks? reconciling natural experiments and semistructural methods. *American Economic Journal: Macroeconomics*, 14(2):96–122.

Commault, J. (2022b). How do persistent earnings affect the response of consumption to transitory shocks *Sciences Po Discussion Paper version*.

Crawley, E. and Theloudis, A. (2024). Income shocks and their transmission into consumption. *arXiv preprint arXiv:2404.12214*.

De Nardi, M., Fella, G., and Paz-Pardo, G. (2020). Nonlinear household earnings dynamics, self-insurance, and welfare. *Journal of the European Economic Association*, 18(2):890–926.

Doris, A., O'Neill, D., and Sweetman, O. (2011). Gmm estimation of the covariance structure of longitudinal data on earnings. *The Stata Journal*, 11(3):439–459.

Dynan, K. E. (1993). How prudent are consumers? *Journal of Political Economy*, 101(6):1104–13.

Dynan, K. E., Skinner, J., and Zeldes, S. P. (2004). Do the rich save more? *Journal of Political Economy*, 112(2):397–444.

Faia, E. and Shabalina, E. (2023). Labour mobility and exposure to risk. *MIMEO*.

Friedman, M. (1957). Introduction to "a theory of the consumption function".

Fuster, A., Kaplan, G., and Zafar, B. (2021). What would you do with $500? spending responses to gains, losses, news, and loans. *The Review of Economic Studies*, 88(4):1760–1795.

Gálvez, J. and Paz-Pardo, G. (2023). Richer earnings dynamics, consumption and portfolio choice over the life cycle. ECB Working Paper 2810, European Central Bank.

Ghosh, A. and Theloudis, A. (2024). Consumption partial insurance in the presence of tail income risk.

Guiso, L., Jappelli, T., and Terlizzese, D. (1992). Earnings uncertainty and precautionary saving. *Journal of Monetary Economics*, 30(2):307–337.

Huggett, M. (1993). The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of Economic Dynamics and Control*, 17(5-6):953–969.

Iao, M. C. T. (2024). Why does unemployment risk reduce the mpc over the business cycle? Job Market Paper, New York University.

Jappelli, T. and Pistaferri, L. (2014). Fiscal policy and mpc heterogeneity. *American Economic Journal: Macroeconomics*, 6(4):107–36.

Jappelli, T. and Pistaferri, L. (2020). Reported mpc and unobserved heterogeneity. *American Economic Journal: Economic Policy*, 12(4):275–97.

Jappelli, T., Savoia, E., and Sciacchetano, A. (2024). Intertemporal mpc and shock size. *Mimeo*.

Kaplan, G., Moll, B., and Violante, G. L. (2018). Monetary policy according to hank. *American Economic Review*, 108(3):697–743.

Kaplan, G. and Violante, G. L. (2010). How much consumption insurance beyond self-insurance? *American Economic Journal: Macroeconomics*, 2(4):53–87.

Kaplan, G. and Violante, G. L. (2021). Propensity to consume in heterogeneous agent models. *Annual Reviews Of Economics (Forthcoming)*.

Kaplan, G., Violante, G. L., and Weidner, J. (2014). The wealthy hand-to-mouth. *Brookings Papers on Economic Activity*, 45(1 (Spring)):77–153.

Kaufmann, K. and Pistaferri, L. (2009). Disentangling insurance and information in intertemporal consumption choices. *American Economic Review: Papers Proceedings*, 99(2):387–92.

Kimball, M. (1990a). Precautionary saving and the marginal propensity to consume. (3403).

Kimball, M. (1990b). Precautionary saving in the small and in the large. *Econometrica*, 58(1):53–73.

Kosar, G. and Melcangi, D. (2024). Subjective uncertainty and the marginal propensity to consume. *Unpublished*.

Krusell, P. and Smith, A. A. (1998). Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy*, 106(5):867–896.

Kueng, L. (2018). Excess Sensitivity of High-Income Consumers. *The Quarterly Journal of Economics*, 133(4):1693–1751.

Laibson, D., Maxted, P., and Moll, B. (2024). Present bias amplifies the household balance-sheet channels of macroeconomic policy. *Quarterly Journal of Economics, Forthcoming*.

Leland, H. (1968). Saving and uncertainty: The precautionary demand for saving. *The Quarterly Journal of Economics*, 82(3):465–473.

Lewis, D. J., Melcangi, D., and Pilossoph, L. (2019). Latent heterogeneity in the marginal propensity to consume. Staff Reports 902, Federal Reserve Bank of New York. Revised July 2021.

Ministero dell'Economia e delle Finanze, Dipartimento delle Finanze (2025). Analisi dei dati delle dichiarazioni irpef - anno d'imposta 2023. https://www1.finanze.gov.it/finanze/analisi_stat/public/v_4_0_0/contenuti/analisi_dati_2023_irpef.pdf. Accessed: 2025-01-15.

Mueller, A. I. (2017). Separations, sorting, and cyclical unemployment. *American Economic Review*, 107(7):2081–2107.

Orchard, J., Ramey, V. A., and Wieland, J. (2025). Micro mpcs and macro counterfactuals: The case of the 2008 rebates. *Quarterly Journal of Economics*. Forthcoming.

Patterson, C. (2023). The matching multiplier and the amplification of recessions. *American Economic Review*, 113(4):982–1012.

Pedroni, M., Singh, S., and Stoltenberg, C. (2022). Advance information and consumption insurance: Evidence from panel data. Tinbergen Institute Discussion Paper TI 2022-032/VI, Tinbergen Institute.

Pratt, J. W. and Zeckhauser, R. (1987). Proper risk aversion. *Econometrica*, 55(1):143–54.

Rouwenhorst, K. G. (1995). Asset pricing implications of equilibrium business cycle models. *Princeton University Press*, page 294–330.

Sala, H. and Trivín, P. (2021). Reported mpc in the presence of debt. *Economics Letters, Elsevier*, 205(C).

Sandmo, A. (1970). The effect of uncertainty on saving decisions. *Review of Economic Studies*, 37(3):353–360.

Sciacchetano, A. (2024). Intertemporal substitution in consumption via subjective expectations: Evidence from a european survey. SSRN Working Paper 4773309, SSRN. Posted: 15 Apr 2024.

Shore, S. H. and Sinai, T. (2010). Commitment, risk, and consumption: Do birds of a feather have bigger nests? *The Review of Economics and Statistics*, 92(2):408–424.

Straub, L. (2019). Consumption, savings, and the distribution of permanent income. *MIMEO*.

Toda, A. A. (2021). Necessity of hyperbolic absolute risk aversion for the concavity of consumption functions. *Journal of Mathematical Economics*, 94:102494.

Wang, T. (2023). Perceived versus calibrated income risks in heterogeneous-agent consumption models. *Bank of Canada Staff Working Paper*, 2023-59.

Zeldes, S. (1989,a). Optimal consumption with stochastic income: Deviations from certainty equivalence. *The Quarterly Journal of Economics*, 104(2):275–298.

Online Appendices for "Intertemporal Insurance and Excess Sensitivity of Affluent Households"

Ettore Savoia
Sveriges Riksbank
March 12, 2026

Contents

1 Introduction **1**

2 An illustrative two-period model **5**

3 The incidence of saving motives and its implications for the MPC **9**
3.1 Distributional Characterization of MPC Rankings 11

4 Stylized Facts **18**
4.1 Empirical strategy 18
4.2 Four Facts on Income Risk Heterogeneity and Household Saving and Consumption Behavior 20
4.3 Evidence from Self-Reported MPCs 24

5 Quantitative Evidence **29**

6 Conclusions **34**

A Mathematical Appendix to Section 2 **1**
A.1 MPC formula derivation 1
A.2 Intertemporal Insurance 3
A.3 Intertemporal Prudence 5
A.4 Discussion on previous results on the effect of income risk on the MPC . 5
A.5 A numerical remark on Kimball (1990a) and Carroll and Kimball (1996) results 8

B Appendix to Section 3 **11**
B.1 Proofs of Distributional Propositions 11
B.2 Different parametrization of the income-risk covariance 12
B.3 Income process with permanent and transitory components 14
B.4 Modeling the negative covariance between current-income and risk 15
B.5 Adding heterogeneity in the borrowing limit 15
B.6 Adding Endogenous Labor Supply 16
B.7 Adding Unemployment Risk 18

B.8 Two-Asset Model 19
B.9 Simpson Paradox 20

C Empirical Appendix to Section 4 **21**
C.1 Descriptive Statistics 21
C.2 BPP estimation Additional results and robustness checks 22
C.3 Details on 2020-2022 Sample 28

D Details on income risk **30**
D.1 Empirical implementation of Step 1 30
D.2 Other Correlates of Income Risk 32
D.3 Income risk decomposition 32

A Mathematical Appendix to Section 2

A.1 MPC formula derivation

Start from the Euler Equation of the two period model:

$$u'(c_0) \;=\; \beta R\,\mathbb{E}_0[u'(c_1)]\,. \tag{25}$$

and recall that:

$$c_0 = x_0 - a_1, \qquad c_1 = Ra_1 + y_1.$$

Also, recall that the period-0 consumption policy is $c_0 = x_0 - a_1^*(x_0)$, where a_1^* solves (25). The current-period MPC is

$$\text{MPC} \;\equiv\; \frac{dc_0}{dx_0}. \tag{26}$$

Derivation

Step 1: Relate MPC to da_1/dx_0. Differentiating $c_0 = x_0 - a_1^*(x_0)$ with respect to x_0 gives

$$\frac{dc_0}{dx_0} \;=\; 1 \;-\; \frac{da_1^*}{dx_0}. \tag{27}$$

Hence MPC $= 1 -$ MPS, where MPS $\equiv \dfrac{da_1^*}{dx_0}$.

Step 2: Implicitly differentiate the Euler equation. Define the Euler residual

$$F(a_1, x_0) \;\equiv\; u'\big(x_0 - a_1\big) \;-\; \beta R\,\mathbb{E}_0\big[u'\big(Ra_1 + y_1\big)\big]\,. \tag{28}$$

At an interior optimum, $F(a_1^*(x_0), x_0) = 0$ by (25). Differentiate $F(a_1^*(x_0), x_0) = 0$ with respect to x_0:

$$\frac{\partial F}{\partial x_0} \;+\; \frac{\partial F}{\partial a_1}\,\frac{da_1^*}{dx_0} \;=\; 0. \tag{29}$$

Compute each partial derivative carefully, using the chain rule and holding the other argument constant:

(i) The x_0-partial:

$$\frac{\partial F}{\partial x_0} = \; u''(x_0 - a_1)\cdot\frac{\partial(x_0 - a_1)}{\partial x_0} = \; u''(c_0)\cdot 1 = \; u''(c_0).$$

(ii) The a_1-partial:

$$\frac{\partial F}{\partial a_1} = \; u''(x_0 - a_1)\cdot\frac{\partial(x_0 - a_1)}{\partial a_1} \;-\; \beta R\,\mathbb{E}_0\Big[u''(Ra_1 + y_1)\cdot\frac{\partial(Ra_1 + y_1)}{\partial a_1}\Big].$$

Since $\dfrac{\partial(x_0 - a_1)}{\partial a_1} = -1$ and $\dfrac{\partial(Ra_1 + y_1)}{\partial a_1} = R$, we get

$$\frac{\partial F}{\partial a_1} = -u''(c_0) - \beta R\, \mathbb{E}_0[u''(c_1) \cdot R] = -u''(c_0) - \beta R^2\, \mathbb{E}_0[u''(c_1)]\,. \tag{30}$$

Step 3: Solve for da_1^*/dx_0. Substitute (30) and $\frac{\partial F}{\partial x_0} = u''(c_0)$ into (29):

$$u''(c_0) + \Big[- u''(c_0) - \beta R^2\, \mathbb{E}_0\big[u''(c_1)\big] \Big] \frac{da_1^*}{dx_0} = 0.$$

Rearrange:

$$\Big[u''(c_0) + \beta R^2\, \mathbb{E}_0\big[u''(c_1)\big] \Big] \frac{da_1^*}{dx_0} = u''(c_0).$$

Hence

$$\frac{da_1^*}{dx_0} = \frac{u''(c_0)}{u''(c_0) + \beta R^2\, \mathbb{E}_0[u''(c_1)]}. \tag{31}$$

Step 4: Obtain the MPC. Using (27) and (31),

$$\begin{aligned} \text{MPC} &= 1 - \frac{da_1^*}{dx_0} = 1 - \frac{u''(c_0)}{u''(c_0) + \beta R^2\, \mathbb{E}_0[u''(c_1)]} \\ &= \frac{\beta R^2\, \mathbb{E}_0[u''(c_1)]}{u''(c_0) + \beta R^2\, \mathbb{E}_0[u''(c_1)]}. \end{aligned}$$

Therefore,

$$\boxed{\text{MPC} = \frac{\beta R^2\, \mathbb{E}_0[u''(c_1)]}{u''(c_0) + \beta R^2\, \mathbb{E}_0[u''(c_1)]}} \tag{32}$$

for the unconstrained (interior) solution.

Signs and Bounds

Under the maintained assumptions, $u''(c_0) < 0$, $u''(c_1) < 0$, and $\mathbb{E}_0[|u''(c_1)|] < \infty$, so $\mathbb{E}_0[u''(c_1)] < 0$. Hence both the numerator and denominator of (32) are negative. Moreover,

$$u''(c_0) + \beta R^2\, \mathbb{E}_0[u''(c_1)] < \beta R^2\, \mathbb{E}_0[u''(c_1)] < 0,$$

so the denominator is more negative than the numerator in absolute value. It follows that

$$0 < \text{MPC} < 1.$$

Borrowing Constraint (Corner Case)

If the borrowing/saving constraint binds at $a_1 = 0$, then the budget constraint implies that $c_0 = x_0$. Differentiating gives $\dfrac{dc_0}{dx_0} = 1$, so

$$\text{MPC} = 1 \quad \text{(by definition at the binding corner)}.$$

Remark: The derivation only uses the Euler equation, the budget identities, and differentiability allowing the interchange of expectation and differentiation (standard regularity). The algebra does not require HARA; the bound MPC $\in (0,1)$ follows from $u''(c) < 0$ and $\mathbb{E}_0[|u''(c_1)|] < \infty$.

A.2 Intertemporal Insurance

This appendix derives the *Intertemporal Insurance* condition presented in proposition 1 of the main text. The condition provides a sufficient criterion for determining whether differences MPC across types are driven by underlying curvature—as opposed to differences in the level of cash-on-hand—even when preferences and interest rates vary across agents.

For each type $i \in \{A, B\}$, the closed-form expression for the MPC in the two-period model is:

$$\text{MPC}_i = \frac{\beta_i R_i^2 \mathbb{E}_0[u''(c_1^i)]}{u''(c_0^i) + \beta_i R_i^2 \mathbb{E}_0[u''(c_1^i)]}.$$

Since $u'' < 0$, define the positive index

$$\mathcal{I}_i^{\mathcal{I}} = \frac{-u''(c_0^i)}{\beta_i R_i^2 \mathbb{E}_0[-u''(c_1^i)]}.$$

Dividing numerator and denominator of the MPC expression by $\beta_i R_i^2 \mathbb{E}_0[u''(c_1^i)]$ yields

$$\text{MPC}_i = \frac{1}{1 + \mathcal{I}_i^{\mathcal{I}}}.$$

Therefore,

$$\text{MPC}_B > \text{MPC}_A \iff \frac{1}{1+\mathcal{I}_B^{\mathcal{I}}} > \frac{1}{1+\mathcal{I}_A^{\mathcal{I}}} \iff \mathcal{I}_B^{\mathcal{I}} < \mathcal{I}_A^{\mathcal{I}}.$$

Special case: identical β and R. If β and R are the same across types, then the ranking condition reduces to

$$\text{MPC}_B > \text{MPC}_A \iff \frac{-u''(c_0^B)}{\mathbb{E}_0[-u''(c_1^B)]} < \frac{-u''(c_0^A)}{\mathbb{E}_0[-u''(c_1^A)]}.$$

The effects of higher β on $\mathcal{I}^{\mathcal{I}}$

Along the equilibrium path, the Euler equation holds:

$$u'(c_0) = \beta R\, \mathbb{E}_0[u'(c_1)].$$

Implicit differentiation with respect to β gives

$$u''(c_0)\frac{dc_0}{d\beta} = R\, \mathbb{E}_0[u'(c_1)] + \beta R\, \mathbb{E}_0\left[u''(c_1)\frac{dc_1}{d\beta}\right].$$

Under standard comparative statics of patience,

$$\frac{dc_0}{d\beta} < 0, \qquad \frac{dc_1}{d\beta} > 0.$$

Consider the total derivative:

$$\frac{d\log \mathcal{I}^{\mathcal{I}}}{d\beta} = \frac{u'''(c_0)}{u''(c_0)}\frac{dc_0}{d\beta} - \frac{1}{\beta} - \frac{\mathbb{E}_0\left[u'''(c_1)\frac{dc_1}{d\beta}\right]}{\mathbb{E}_0[u''(c_1)]}.$$

Using $u'' < 0$, $u''' > 0$, $\frac{dc_0}{d\beta} < 0$, and $\frac{dc_1}{d\beta} > 0$, the first and third terms are positive, while the middle term is negative. Hence, without further restrictions, the sign of $\frac{d\log \mathcal{I}^{\mathcal{I}}}{d\beta}$ is in general ambiguous.

CRRA illustration

Assume CRRA preferences $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$ with $\gamma > 0$. Then

$$u'(c) = c^{-\gamma}, \qquad u''(c) = -\gamma c^{-\gamma-1}.$$

The Euler equation implies

$$c_0^{-\gamma} = \beta R\, \mathbb{E}_0[c_1^{-\gamma}].$$

The intertemporal insurance index becomes

$$\mathcal{I}^{\mathcal{I}} = \frac{-u''(c_0)}{\beta R^2\, \mathbb{E}_0[-u''(c_1)]} = \frac{c_0^{-\gamma-1}}{\beta R^2\, \mathbb{E}_0[c_1^{-\gamma-1}]}.$$

Certainty case. If uncertainty vanishes, $\mathbb{E}_0[c_1^{-\gamma-1}] = c_1^{-\gamma-1}$, and the Euler equation gives

$$\frac{c_1}{c_0} = (\beta R)^{1/\gamma}.$$

Substituting,

$$\mathcal{I}^{\mathcal{I}} = \frac{1}{\beta R^2}\left(\frac{c_1}{c_0}\right)^{\gamma+1} = \frac{(\beta R)^{(\gamma+1)/\gamma}}{\beta R^2} = \beta^{1/\gamma} R^{1/\gamma - 1}.$$

Therefore,

$$\frac{\partial \log \mathcal{I}^{\mathcal{I}}}{\partial \beta} = \frac{1}{\gamma\beta} > 0, \qquad \Rightarrow \qquad \frac{\partial \mathcal{I}^{\mathcal{I}}}{\partial \beta} = \frac{\mathcal{I}^{\mathcal{I}}}{\gamma\beta} > 0.$$

Hence, in the CRRA certainty benchmark, higher patience increases $\mathcal{I}^{\mathcal{I}}$ and lowers the MPC.

□

A.3 Intertemporal Prudence

For any type i in the Euler-equation region,

$$MPC_i = \frac{1}{1+\mathcal{I}_i^{\mathcal{I}}}, \qquad \mathcal{I}_i^{\mathcal{I}} = \frac{-u_i''(c_0^i)}{\beta_i R_i^2\, \mathbb{E}_0[-u_i''(c_1^i)]}.$$

To obtain a local fixed-cash characterization, take a first-order Taylor expansion of $-u_i''(c_1^i)$ around c_0^i:

$$-u_i''(c_1^i) \approx -u_i''(c_0^i) - u_i'''(c_0^i)(c_1^i - c_0^i).$$

Taking conditional expectations gives

$$\mathbb{E}_0[-u_i''(c_1^i)] \approx -u_i''(c_0^i) - u_i'''(c_0^i)(\mathbb{E}_0[c_1^i] - c_0^i).$$

Factorizing $-u_i''(c_0^i)$:

$$\mathbb{E}_0[-u_i''(c_1^i)] \approx -u_i''(c_0^i)\Big[1 + \underbrace{\left(-\frac{u_i'''(c_0^i)}{u_i''(c_0^i)}(c_0^i - \mathbb{E}_0[c_1^i])\right)}_{\mathcal{P}_i^{\mathcal{I}}}\Big].$$

Hence,

$$\mathcal{I}_i^{\mathcal{I}} \approx \frac{-u_i''(c_0^i)}{\beta_i R_i^2\,[-u_i''(c_0^i)(1+\mathcal{P}_i^{\mathcal{I}})]} = \frac{1}{\beta_i R_i^2(1+\mathcal{P}_i^{\mathcal{I}})}.$$

Substituting into the MPC formula,

$$MPC_i = \frac{1}{1+\mathcal{I}_i^{\mathcal{I}}} \approx \frac{\beta_i R_i^2(1+\mathcal{P}_i^{\mathcal{I}})}{1+\beta_i R_i^2(1+\mathcal{P}_i^{\mathcal{I}})}.$$

Since the function

$$\Lambda \mapsto \frac{\Lambda}{1+\Lambda}$$

is strictly increasing for $\Lambda > 0$, it follows that at any fixed $\bar{x}_0$,

$$MPC_B > MPC_A \quad \Longleftrightarrow \quad \beta_B R_B^2(1+\mathcal{P}_B^{\mathcal{I}}) > \beta_A R_A^2(1+\mathcal{P}_A^{\mathcal{I}}).$$

If $\beta_A R_A^2 = \beta_B R_B^2$, the ranking reduces to

$$MPC_B > MPC_A \quad \Longleftrightarrow \quad \mathcal{P}_B^{\mathcal{I}} > \mathcal{P}_A^{\mathcal{I}}.$$

□

A.4 Discussion on previous results on the effect of income risk on the MPC

As mentioned in the main text, my paper adopts a different perspective from classic studies such as Kimball (1990b) and Carroll and Kimball (1996). Those papers characterize how income risk changes the consumption policy function itself, in particular by increas-

ing its concavity, and evaluate the effects of uncertainty within a given type. My analysis instead compares effective allocations across types and asks how consumption behavior differs once savings choices adjust endogenously to different levels of risk. To illustrate this distinction transparently, below, I define and introduce compensating and equivalent precautionary premia within the two-period model studied in the main text, and later I quantify numerically the effects of these approach in the infinite-horizon model.

1. Kimball (1990a,b) study "risk-compensated" allocations that restore the certainty-equivalent consumption allocation in the presence of risk using a compensating precautionary premium.
2. Carroll and Kimball (1996) characterize the optimal risky consumption function and relate MPC implications to its curvature at a given wealth (cash-on-hand) state; they also discuss an "equivalent premium" device to restore certainty-equivalent allocations at a fixed state.

Revisiting Kimball (1990a,b). If we want to revisit the analysis in the sense of Kimball, we should ask what allocation (and implied MPC) would a risk-exposed agent display in the absence of income uncertainty. To answer this, let us assume we can use a *compensating precautionary premium* ϕ^*. That is, a lump-sum transfer that offsets risk and restores the certainty-equivalent consumption allocation. Although Kimball's framework is slightly more complicated than the simple model presented above, Toda (2021) uses his two-period setup, and delivers exactly the same MPC formula as equation (7), allowing for a direct comparison. For comparability, as in Kimball, I set $\beta R = 1$ and $R = 1$, so that the certainty-equivalent MPC is 1/2.

The "risk-compensated" Euler equation (5) then reads:

$$u'\big(c_0^{\text{ce}}\big) \;=\; \mathbb{E}_0\big[u'(c_1 + \phi^*)\big].$$

Importantly, this counterfactual adjusts current consumption so as to restore the Euler condition, while the risky future consumption profile remains unchanged. Hence the future-curvature term entering the MPC formula continues to be evaluated at $E_0[u''(c_1)]$. The resulting transfer makes the uncertain consumer able to attain the (higher) certainty-equivalent current consumption level. Since u' is strictly decreasing, adding $\phi^* > 0$ to future consumption lowers expected marginal utility to the certainty-equivalent level. To understand the logic behind this exercise, Figure A.4.1a illustrates the effect on marginal utility and, via curvature, on the MPC formula (7). Under this risk-compensated comparison, the current-curvature term $u''(c_0^*)$ is replaced by $u''(c_0^{ce})$ in equation 7, while future curvature remains evaluated at the risky c_1. Hence, equation 7 becomes:

$$\overline{\text{MPC}} = \frac{\mathbb{E}_0[u''(c_1)]}{u''(c_0^{\text{ce}}) + \mathbb{E}_0[u''(c_1)]}\,.$$

Since the certainty-equivalent allocation implies higher current consumption, standard curvature (e.g., CRRA) delivers lower absolute curvature at c_0^{ce}, i.e. $|u''(c_0^{\text{ce}})|$ is smaller. Under risk, $|\mathbb{E}_0[u''(c_1)]|$ is larger due to the convexity of $-u''(\cdot)$ under standard preferences. Hence $|u''(c_0^{\text{ce}})| < |\mathbb{E}_0[u''(c_1)]|$, and so $\overline{\text{MPC}} > 1/2$.

This conclusion is specific to the risk-compensated exercise, which fixes current curvature at c_0^{ce} by construction. Along the optimal policy under risk both c_0 and the distribution of c_1 adjust endogenously. In particular, precautionary saving implies $c_0 \leq E_0[c_1]$ when $u''' > 0$ and $\beta R = 1$. From equation (7),

$$MPC = \frac{1}{1+\mathcal{I}}, \qquad \mathcal{I} = \frac{-u''(c_0)}{E_0[-u''(c_1)]},$$

so the effect of income risk operates through the intertemporal insurance term $\mathcal{I}$. An increase in precautionary saving raises the degree of intertemporal insurance and therefore lowers the MPC along the optimal consumption function.

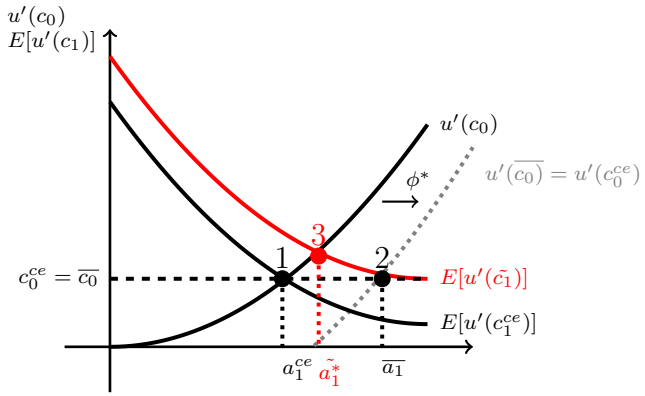
Figure A.4.1c illustrates the compensating precautionary premium in action using consumption and MPC–cash-on-hand schedules. The black lines represent the certainty-equivalent case, where consumption is linear and the MPC is constant at 1/2. The red curves show the (uncompensated) schedules under uncertainty: at the borrowing constraint the MPC equals 1, and away from the constraint the interior MPC is given by (7). The pink "compensated" curve restores the certainty-equivalent allocation, which mechanically raises the MPC relative to the uncompensated risky case. Hence, without the compensating precautionary premium, the MPC under uncertainty lies above the certainty-equivalent MPC only at the constraint.

Revisiting Carroll and Kimball (1996). The same idea applies here, but at a fixed wealth (cash-on-hand) state across environments. The tool here is an *equivalent precautionary premium*. That is, a lump-sum transfer that offsets risk and restores the level of wealth the household would have achieved had uncertainty been absent. In our two-period notation, a convenient representation of this counterfactual is:

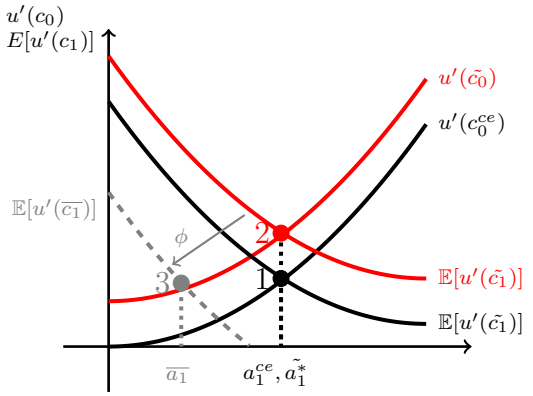
$$u'\big(c_0 - \phi\big) = \mathbb{E}_0\big[u'(c_1)\big].$$

Equivalently, one can represent the same compensation as a shift in the expected marginal-utility term; we use the left-hand representation to emphasize that the device wedges current marginal utility to restore the certainty-equivalent allocation at a fixed state. Figure A.4.1b illustrates the marginal-utility logic and its implication for the curvature term entering (7). As in the previous exercise, the compensated comparison replaces the current-curvature term $u''(c_0^*)$ by $u''(c_0^{ce})$, generating the same counterfactual force toward MPC $> 1/2$.

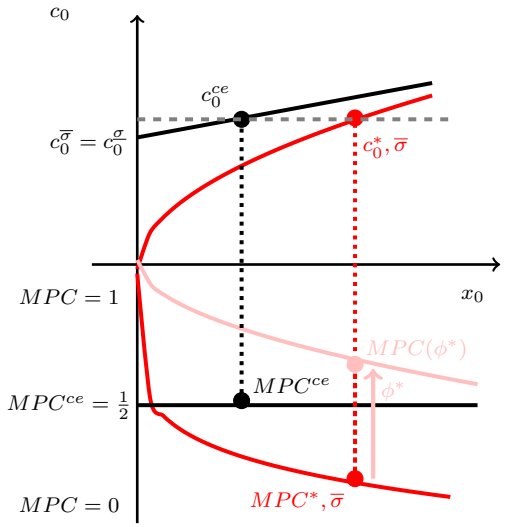
This analysis clarifies a subtle point that is commonly confused. In the two-period model with $\beta R = 1$, it is not the precautionary motive *per se* that raises the MPC at a given state, but rather the counterfactual compensation devices (ϕ^* and ϕ) used to restore certainty-equivalent allocations and quantify the cost of risk. Below I now show that such counterfactual exercises applied to infinite-horizon incomplete-markets models with household heterogeneity, also generate small quantitative differences in MPCs.



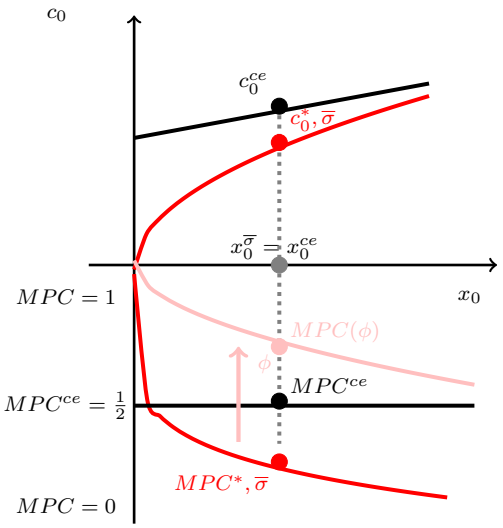
(a) Kimball's compensating premium ϕ^* restores the certainty-equivalent current consumption level c_0^{ce} under risk by shifting the marginal-utility condition. The labeled points depict (1) the certainty-equivalent allocation, (2) the higher saving required to attain c_0^{ce} under risk, and (3) the risky optimum.



(b) Carroll–Kimball's "equivalent premium" ϕ restores the certainty-equivalent allocation at a fixed state. Point 1 denotes the certainty-equivalent saving level a_1^{ce}. Evaluated on the risky consumption schedule this corresponds to a lower level of current consumption and therefore higher marginal utility (point 2). Point 3 denotes the risky optimum $(\tilde{c}_0^*, \tilde{a}_1^*)$.



(c) The red curves show the (uncompensated) consumption function and MPC under risk. At any given x_0 in $\mathcal{X}$, the risky schedule implies a lower local MPC than under certainty. By contrast, the pink "compensated" curve restores the certainty-equivalent allocation and mechanically raises the MPC relative to the uncompensated risky case.



(d) The red curves show the (uncompensated) risky schedule at a fixed cash-on-hand state. The pink "compensated" curve restores the certainty-equivalent allocation at that state and mechanically raises the MPC relative to the uncompensated risky case.

Figure A.4.1: Euler-equation representation and MPC–cash-on-hand curves illustrating the role of compensating premia in Kimball (1990b) (Panels A.4.1a and A.4.1c) and Carroll and Kimball (1996) (Panels A.4.1b and A.4.1d).

A.5 A numerical remark on Kimball (1990a) and Carroll and Kimball (1996) results

The left panels show a scenario where households, despite facing different income risk levels, are counterfactually imposed to consume the same amount—verifying the mechanism behind Kimball (1990a,b). The right panels repeat this exercise but keep assets constant instead, validating the mechanism behind Carroll and Kimball (1996). These

solutions are obtained by adjusting the discount factor at each risk level to ensure the consumption or asset target is met. The target is merely illustrative and can be any value.[27]

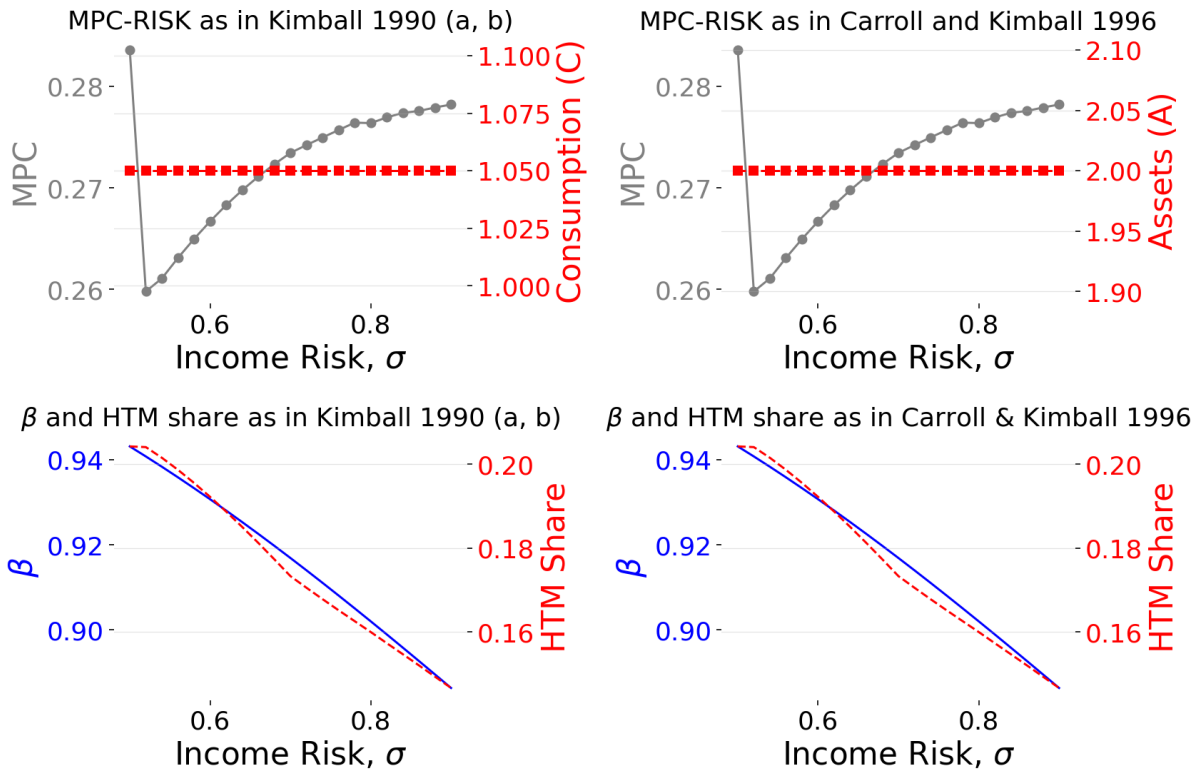


Figure A.5.1: MPC, assets, and risk: counterfactual relationships

The figure illustrates a simple trade-off. As income risk σ increases, the precautionary motive strengthens and the intertemporal insurance channel becomes active. Within each riskier income process, households front-load less and therefore the stationary asset distribution is shifted away from the borrowing limit, so the HTM (constrained or near-constrained) share tends to fall. Along our calibration, since we imposed the same average consumption and wealth, however, the discount factor β must decline with σ, which works in the opposite direction by increasing impatience and pulling households closer to the constraint. In the simulations shown, the precautionary-saving channel dominates, hence the HTM share is mildly decreasing even though β is falling.

The implied effect of risk on the average MPC is nevertheless small. The reason is that the two forces largely offset in the cross-section: higher σ pushes mass toward higher wealth where MPCs are lower, while lower β shifts mass back toward lower wealth where MPCs are higher. Because the average MPC is a probability-weighted average of individual MPCs, it changes substantially only if the distribution moves a lot across regions with very different local MPCs. Here that reallocation is limited (most mass remains in a region where MPC varies smoothly), so the MPC–risk slope is positive but quantitatively modest.

[27]The risk values range from 0.50 to 0.90 for convenience, as alternative values do not affect the results. However, imposing zero risk is irrelevant since in this model zero assets are implying, making it impossible to achieve any positive target.

Overall, this evidence confirms that although the Kimball (1990a,b); Carroll and Kimball (1996) framework is excellent in quantifying risk costs in intertemporal allocation, it is less suitable to assess how income risk heterogeneity shapes MPC heterogeneity.

B Appendix to Section 3

This appendix focuses on the covariance between income risk and income levels, as this constitutes the main application of the paper. Nonetheless, similar patterns emerge under alternative parameterizations of the baseline specification featuring heterogeneity in the relative risk aversion and in the discount factor, as documented in the main text. These additional results are not reported here for brevity, but they are available upon request. As reported in the main text, Figures in this appendix also report results for the positive income incidence and for the implied MPC-permanent income scalar ψ.

B.1 Proofs of Distributional Propositions

Throughout, assume:

(i) The asset space is $a \in [0, \infty)$, (ii) Each type-i stationary distribution F_i admits a density where needed, (iii) MPC schedules $m_i(a)$ are measurable and weakly decreasing in a, (iv) Differentiability conditions stated in the propositions hold pointwise.

All expectations are taken with respect to the stationary distributions.

B.1.1 Proof of Proposition 4

For type i,

$$\bar{m}_i = D_{0,i} + (1 - D_{0,i})\bar{m}_i^E.$$

Consider two types i and j.

Then

$$\bar{m}_i - \bar{m}_j = (D_{0,i} - D_{0,j}) + [(1 - D_{0,i})\bar{m}_i^E - (1 - D_{0,j})\bar{m}_j^E].$$

Rearranging,

$$\bar{m}_i - \bar{m}_j = (D_{0,i} - D_{0,j}) - [(1 - D_{0,j})\bar{m}_j^E - (1 - D_{0,i})\bar{m}_i^E].$$

Hence $\bar{m}_i < \bar{m}_j$ if and only if

$$(D_{0,j} - D_{0,i}) > (1 - D_{0,i})\bar{m}_i^E - (1 - D_{0,j})\bar{m}_j^E.$$

This establishes the stated condition.

Note that this comparison concerns stationary averages only and does not impose any pointwise ranking of $m_i(a)$ and $m_j(a)$. □

B.1.2 Proof of Proposition 5

Define the aggregate MPC schedule

$$M(a) = \sum_{i \in \mathcal{T}} \pi_i(a) m_i(a),$$

where $\pi_i(a) = \Pr(i \mid a)$.
Assume differentiability in a.
Then

$$M'(a) = \sum_i \pi_i'(a) m_i(a) + \sum_i \pi_i(a) m_i'(a).$$

Rearranging,

$$M'(a) = \underbrace{\sum_i \pi_i(a) m_i'(a)}_{\text{within-type effect}} + \underbrace{\sum_i \pi_i'(a) m_i(a)}_{\text{sorting effect}}.$$

This establishes the first decomposition.

Covariance Representation

Using $\sum_i \pi_i(a) = 1$, we have $\sum_i \pi_i'(a) = 0$.
Define expectation conditional on a as:

$$\mathbb{E}_a[X_i] = \sum_i \pi_i(a) X_i.$$

Then

$$\sum_i \pi_i'(a) m_i(a) = \mathrm{Cov}_a\left(m_i(a), \pi_i'(a)/\pi_i(a)\right).$$

Equivalently, writing directly,

$$M'(a) = \mathbb{E}_a[m_i'(a)] + \mathrm{Cov}_a\left(m_i(a), \partial_a \log \pi_i(a)\right).$$

Thus

$$M'(a) > 0$$

if and only if the sorting covariance term dominates the average within-type decline.
This completes the proof. □

B.2 Different parametrization of the income-risk covariance

Figure B.2.1 illustrates how the slope of the income–risk covariance shapes the cross-sectional MPC pattern. The key insight is that the faster income falls as risk rises, the more risk becomes concentrated among households with limited capacity to save. This flattens the asset–risk profile and steepens the MPC–cash profile, since high-income households operating on less concave consumption functions target less wealth. Notably, the MPC–risk curve is unaffected. As long as households have some capacity to save (i.e. are on their Euler equation), they smooth consumption as much as possible, which lowers the MPC regardless of the shape of the covariance.

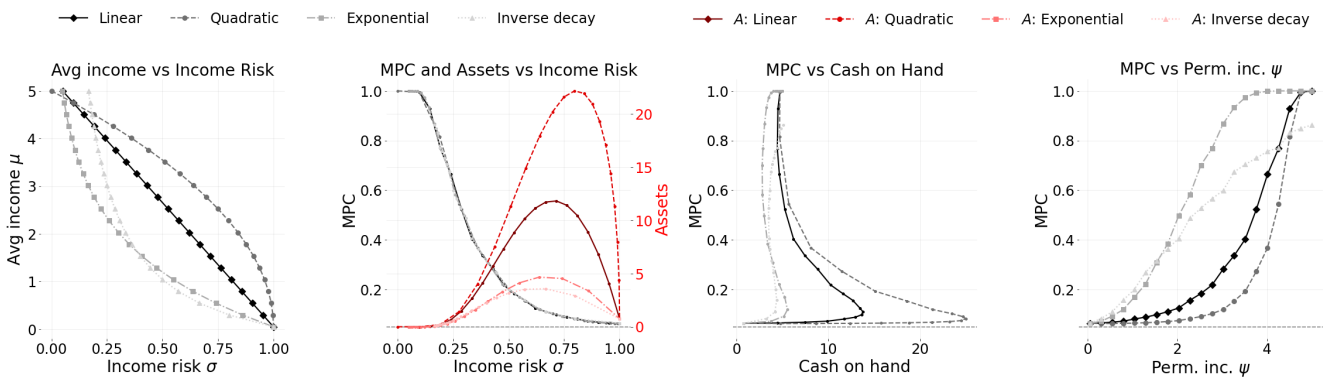


Figure B.2.1: MPC, assets, risk, cash-on-hand, and income relationships for different shapes of the covariance income-risk

Notes: The figure displays how the curvature of the income–risk covariance affects the relationships among MPC, assets, cash-on-hand and permanent income. The four panels compare linear, quadratic, exponential and inverse–decay shapes of the covariance.

B.3 Income process with permanent and transitory components

In this section, I consider a version of the baseline model presented in Section 2, where the only difference is that now a household of type j has the same income process as in Section 5. Specifically, idiosyncratic income is the exponential of a persistent and a transitory component:

$$\log y_t = z_t + \varepsilon_t, \tag{33}$$

$$z_t = \rho z_{t-1} + \eta_t, \tag{34}$$

The covariance is modeled through the time-invariant scaling of the income process, $\psi_{i,y} \times y_{i,t}$. Figure B.3.1 considers the case whereas the standard deviation of the persistent component raises, (σ_η), the scalar $\psi_{i,y}$ decreases. While Figure B.3.2 considers the case whereas the standard deviation of the transitory component raises, (σ_ε), the scalar $\psi_{i,y}$ decreases.

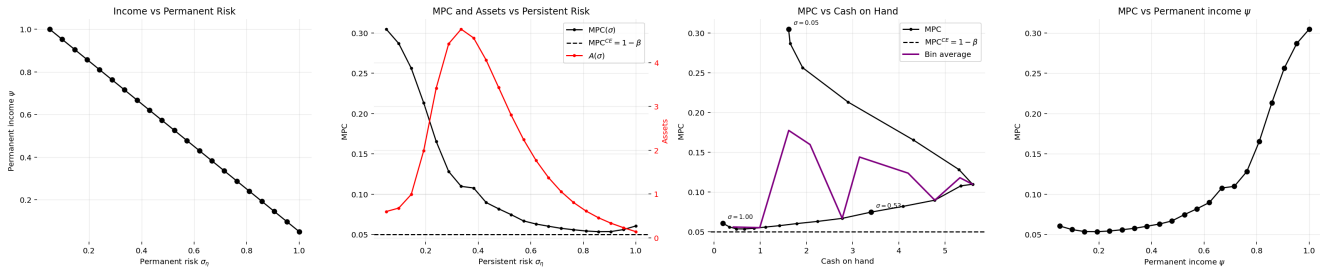


Figure B.3.1: MPC, Asset, risk and Cash-on-hand relationships under Negative Incidence Permanent Income-Permanent Risk

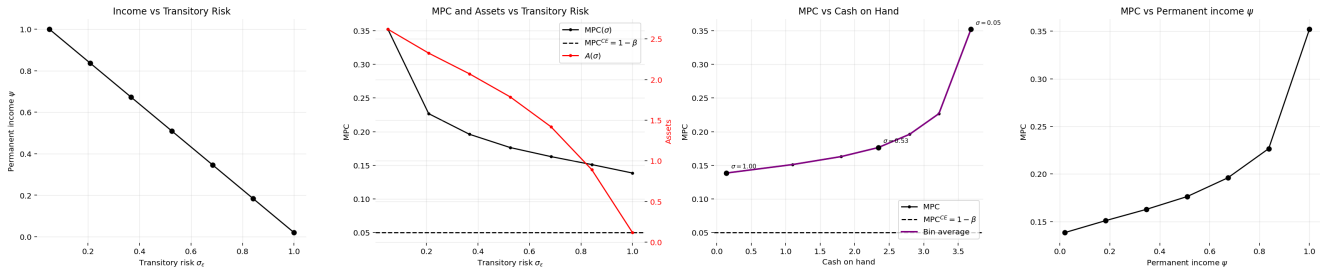


Figure B.3.2: MPC, Asset, risk and Cash-on-hand relationships under Negative Incidence Permanent Income-Transitory Risk

B.4 Modeling the negative covariance between current-income and risk

In this section, I model the negative relationship between income risk and income level by directly specifying $y = \exp(s - \mu(\sigma))$, where s is the idiosyncratic state and $\mu(\sigma) = \max(s) + \omega \cdot \sigma^k$ with $k, \omega \geq 1$ controlling the shape of the covariance. The baseline assumes $\kappa = 2, \omega = 1$. While this formulation is algebraically equivalent to a type-specific scaling of income, it interprets the restriction as operating through current income levels rather than through an explicit permanent-income scalar ψ. Figure B.4.1 reports the MPC-risk, asset-risk, MPC-cash and MPC-average income for different slopes of the new income–risk covariance.

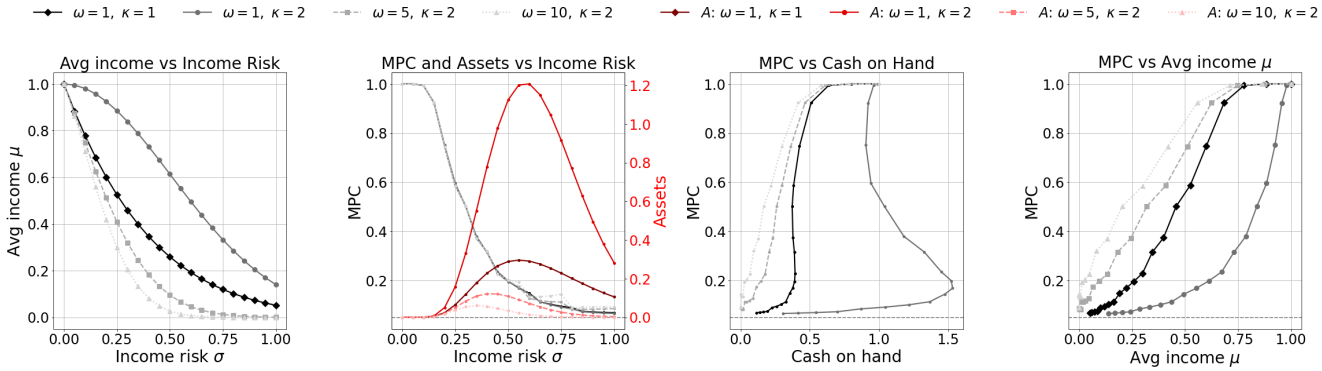


Figure B.4.1: MPC, asset, risk, cash-on-hand and income relationships for different shapes of the covariance income-risk

Notes: The figure shows the MPC-risk, asset-risk, MPC-cash and MPC-average income for different slopes of the income–risk covariance. Recall that $y = \exp\{s - \mu(\sigma)\}$ with $\mu(\sigma) = \max(s) + \omega\,\sigma^{\kappa}$ and $\kappa, \omega \geq 1$. Since $\mu'(\sigma) = \omega\kappa\,\sigma^{\kappa-1} > 0$, larger ω or κ makes mean income fall faster as risk σ rises, which is what the first column shows (black: $\omega = 1, \kappa = 1$; progressively lighter grays: $\omega = 1, \kappa = 2$, $\omega = 5, \kappa = 2$, and finally $\omega = 10, \kappa = 2$). Main takeaways: (i) the steeper the covariance, the flatter the asset–risk curve; (ii) the steeper the covariance, the more increasing the MPC–cash and MPC–income curves become.

B.5 Adding heterogeneity in the borrowing limit

A potential concern is that households with higher income and lower risk may have greater access to credit markets. To address this, I allow the borrowing limit to vary with income risk, such that households facing higher income risk are characterized by both lower income levels and tighter borrowing constraints. The dynamic results remain unaffected. This confirms, in a cross-sectional setting, the well-known result emphasized by Carroll et al. (2021): adding liquidity constraints on top of precautionary motives does not generate economically meaningful changes in the MPC.

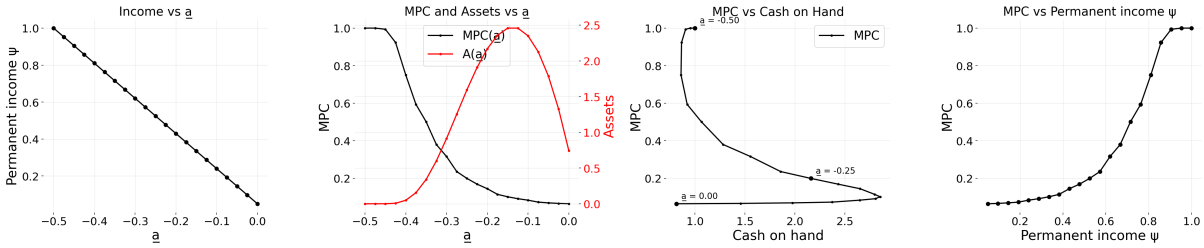


Figure B.5.1: MPC, assets, income risk, cash-on-hand, and permanent income relationships under heterogeneous borrowing constraints, income risk, and income levels

B.6 Adding Endogenous Labor Supply

Compared to the baseline standard incomplete-markets income model presented in equation 10, this section extends the framework to incorporate endogenous labor supply. The model reproduces the main comparative statics presented in the text, confirming that the central results remain robust.

$$V_t(y_{i,t}, a_{i,t-1}) = \max_{c_{i,t}, n_{i,t}, a_{i,t}} \left\{ \frac{c_{i,t}^{1-\gamma}}{1-\gamma} - \varphi \frac{n_{i,t}^{1+\nu}}{1+\nu} + \beta \, \mathbb{E}_t[V_{t+1}(y_{i,t+1}, a_{i,t})] \right\}$$

$$\text{s.t.} \qquad c_{i,t} + a_{i,t} = (1+r)a_{i,t-1} + \psi \cdot y_{i,t}\, n \qquad \text{and} \qquad a \geq 0.$$

$$logy_{i,t} = \rho \log y_{i,t-1} + \varepsilon_{i,t}, \tag{35}$$

where $\varepsilon_{i,t}$ is an i.i.d. normal process with mean zero and standard deviation σ_ε.

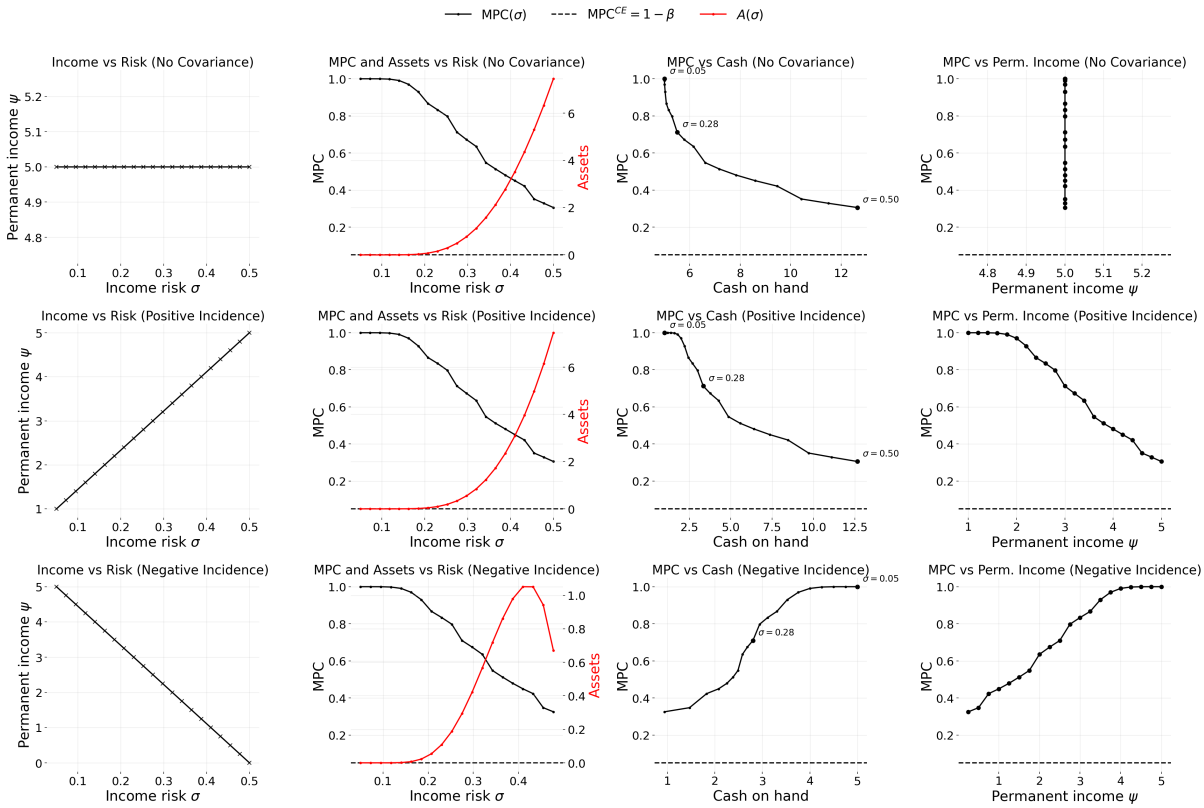


Figure B.6.1: MPC, assets, risk, cash-on-hand, and permanent-income relationships with endogenous labor supply

Notes: The figure contrasts three income–risk covariance structures: none (top), positive (middle), and negative (bottom). For each case, the second column plots steady-state average MPC (black) and assets (red) across risk–income combinations; the third shows the implied MPC–cash-on-hand relationship; and the fourth shows the implied MPC–permanent-income profile, which mirrors the MPC–risk relation.

B.7 Adding Unemployment Risk

In this section, I show that the main intuition continues to hold even when risk is represented by separation risk in a simple model with labor force participation and unemployment risk. The model now is written so that employment status matters for income. Employed workers earn a wage, while unemployed workers receive benefits.

Employment status evolves exogenously.[28] An employed worker may lose her job with probability σ. An unemployed worker may find a job with probability f. Let Π^s denote the Markov transition matrix associated with the employment state s_i. These transition probabilities are time-invariant parameters in the stationary environment considered here.

The Bellman equation is

$$V_t(s_{i,t}, y_{i,t}, a_{i,t-1}) = \max_{c_{i,t}, a_{i,t}} \left\{ \frac{c_{i,t}^{1-\gamma}}{1-\gamma} + \beta \, \mathbb{E}_t[\, V_{i,t+1}(s_{i,t+1}, y_{i,t+1}, a_{i,t})] \right\}$$

$$\text{s.t.} \qquad c_{i,t} + a_{i,t} = (1+r)a_{i,t-1} + \psi \cdot y_{i,t}, \qquad a_t \geq \underline{a}.$$

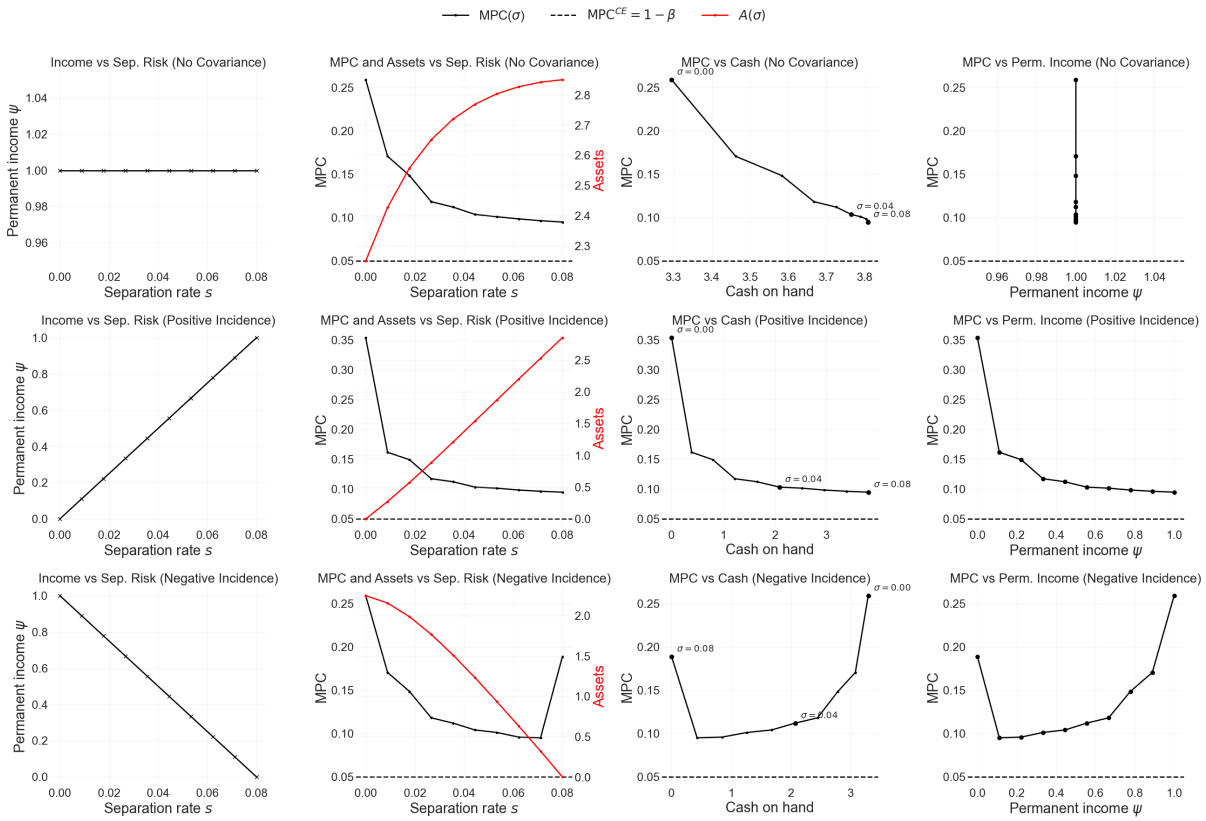


Figure B.7.1: MPC, asset, separation risk, cash-on-hand and permanent income relationships with labor participation and unemployment risk

Notes: The figure contrasts three income–risk covariance structures: none (top), positive (middle), and negative (bottom). For each case, the second column plots steady-state average MPC (black) and assets (red) across separation risk–income combinations; the third shows the implied MPC–cash-on-hand relationship; and the fourth shows the implied MPC–permanent-income profile, which mirrors the MPC–risk relation.

[28] In Challe and Savoia (2026) we build on these insights and show that the result extends to a HANK model with endogenous search and matching.

B.8 Two-Asset Model

In this section I instead show the mechanism in a standard two-asset model with adjustment costs as in Kaplan et al. (2018).

The household problem is characterized by the Bellman equation

$$V_{i,t}(e_{i,t}, b_{i,t-1}, a_{i,t-1}) = \max_{c_{i,t},\, b_{i,t},\, a_{i,t}} \left\{ \frac{c_{i,t}^{1-\gamma}}{1-\gamma} + \beta\, \mathbb{E}_t[V_{i,t+1}(e_{i,t+1}, b_{i,t}, a_{i,t}) \mid e_{i,t}] \right\}$$

$$\text{s.t.} \qquad c_{i,t} + a_{i,t} + b_{i,t} = z_t(e_{i,t}) + (1+r^a)a_{i,t-1} + (1+r^b)b_{i,t-1} - \Psi(a_{i,t}, a_{i,t-1}),$$

$$a_{i,t} \geq \underline{a}, \qquad b_{i,t} \geq \underline{b}.$$

where $z_t(e_{i,t})$ denotes labor income. The adjustment cost function is specified as

$$\Psi(a_{i,t}, a_{i,t-1}) = \frac{\chi_1}{\chi_2} \left| \frac{a_{i,t} - (1+r^a)a_{i,t-1}}{(1+r^a)a_{i,t-1} + \chi_0} \right|^{\chi_2} \left[(1+r^a)a_{i,t-1} + \chi_0\right].$$

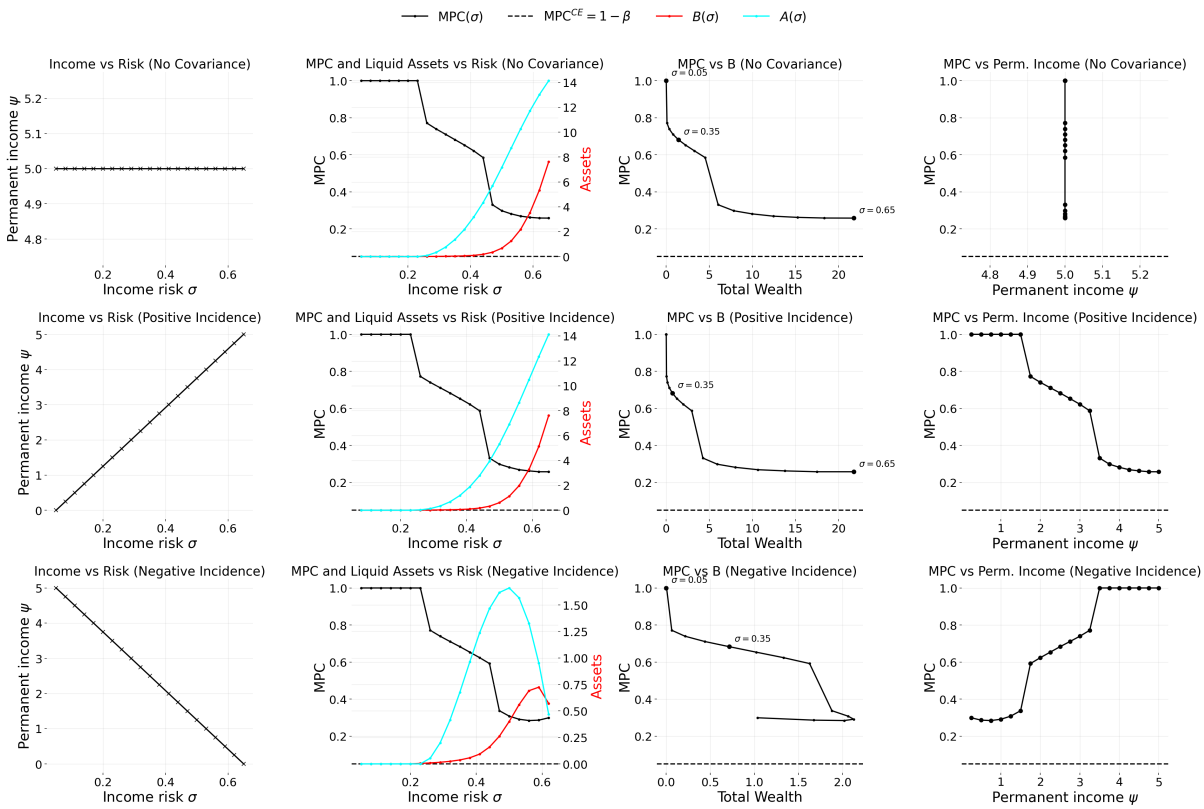


Figure B.8.1: MPC, Assets, Risk and Permanent Income relationships in a two-asset model

Notes: The figure contrasts three income–risk covariance structures: none (top), positive (middle), and negative (bottom). In each row, the first column plots permanent income against income risk. The second column plots steady-state average MPC (black), liquid assets (red), and illiquid assets (cyan) as functions of income risk across risk–income combinations. The third column shows the implied MPC–total-wealth relationship, and the fourth column reports the corresponding MPC–permanent-income profile.

B.9 Simpson Paradox

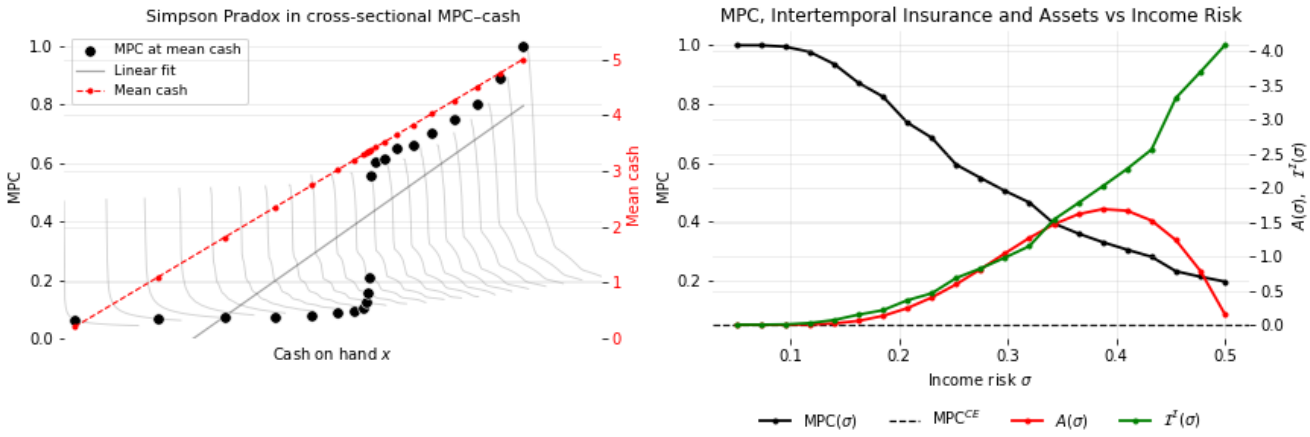


Figure B.9.1: Simpson Paradox in the Standard Incomplete Market Model with Heterogeneous Risk

Notes: Left panel. Each thin gray curve plots the cross–sectional relation between the MPC and cash-on-hand within a given income–risk process. Black dots mark the MPC evaluated at the group's mean cash level, and the solid black line is the least–squares fit through those representative points. The red dashed line (right axis) reports the mean cash for each group. The figure illustrates a Simpson-type pattern: although MPC typically falls with cash *within* a curve, shifts in risk move the representative points along/outside the curves, generating an *across-group* relation that can differ in slope from the within-curve relation. *Right panel.* The black line shows the average MPC as a function of income risk σ; the dashed horizontal line is the certainty–equivalent benchmark $\text{MPC}^{CE} = 1 - \beta$. The red line (right axis) reports average assets $A(\sigma)$, while the green line reports the intertemporal–insurance coefficient $\mathcal{I}^{\mathcal{I}}(\sigma) = (1 - \text{MPC})/\text{MPC}$. As risk rises, the MPC declines and approaches the certainty–equivalent value at high risk, while intertemporal insurance increases; assets rise with risk at first and may decline at very high risk.

C Empirical Appendix to Section 4

This appendix provides additional details on the data used in the main text and presents supplementary exercises that support the main stylized findings.

C.1 Descriptive Statistics

Table C.1.1 reports descriptive statistics for the full sample covering the period 1998–2022 used in the main text to document Facts 1–8. Statistics are reported for the entire sample, as well as separately for low-risk households (below the median of the risk (CV) distribution) and high-risk households (above the median of the risk (CV) distribution).

Table C.1.1: Descriptive Statistics By risk group, SHIW Sample 1998-2022.

	Full sample		high-risk		low-risk	
	Mean	SD	Mean	SD	Mean	SD
Demographics						
Male	0.69	0.46	0.66	0.47	0.72	0.45
Age	45.13	6.78	42.72	8.07	47.54	3.90
Married	0.74	0.44	0.64	0.48	0.85	0.36
Years of education	1.63	0.48	1.62	0.49	1.63	0.48
Resident in the South	0.28	0.45	0.33	0.47	0.22	0.42
Size of the family	3.17	1.26	2.81	1.31	3.52	1.08
Number of labor income earners	1.68	0.66	1.49	0.59	1.87	0.67
City size $\geq$ 500,000	0.08	0.28	0.09	0.28	0.08	0.27
Financials (×1,000)						
Disposable labor income	39.20	20.15	31.35	16.16	47.04	20.67
Net Cash-on-hand	30.56	86.52	20.03	45.37	41.07	112.62
Net liquid wealth	27.03	85.93	17.21	44.95	36.83	112.02
Net illiquid wealth	173.48	329.22	119.99	183.96	226.89	420.88
Non-durable consumption	26.86	13.51	22.42	9.81	31.29	15.16
Liquid Asset-to-income ratio	0.58	1.46	0.50	1.32	0.67	1.58
Average Propensity to Consume	0.74	0.27	0.79	0.31	0.70	0.21
Income Risk						
Coefficient of Variation (CV)	0.05	0.04	0.06	0.04	0.03	0.01
Number of observations	9,779		4,886		4,893	

Notes: The table reports the sample statistics for the main sample used Facts 1-4 of the main text. Households are classified as *low-risk* (*high-risk*) if their income risk measure lies below (above) the median of the CV distribution as presented in the main text. Differences in income level, risk (CV) and financial holdings are statistically different from zero ($p < 1\%$). Following Jappelli and Pistaferri (2014, 2020), net cash-on-hand is defined as net liquid wealth plus one month of disposable income (i.e., annual disposable income divided by 12).

C.2 BPP estimation Additional results and robustness checks

Variance-Autocovariance Matrix of income and consumption growth Table C.2.1 provides estimated covariances between detrended log consumption growth and current and future detrended log income growth.

Table C.2.1: The Variance-Autocovariance Matrix Of Income And Consumption Growth

	$\Delta \ln(y_{i,t})$	$\Delta \ln(y_{i,t+1})$	$\Delta \ln(y_{i,t+2})$	$\Delta \ln(y_{i,t+3})$
$\text{cov}(\Delta \ln(y_{i,t}), \cdot)$	0.074***	-0.025***	-0.004**	-0.001
	(0.004)	(0.002)	(0.001)	(0.002)
$\text{cov}(\Delta \ln(c_{i,t}), \cdot)$	0.038***	-0.012***	-0.001	0.003
	(0.003)	(0.002)	(0.002)	(0.002)
	$\Delta \ln(c_{i,t})$	$\Delta \ln(c_{i,t+1})$	$\Delta \ln(c_{i,t+2})$	$\Delta \ln(c_{i,t+3})$
$\text{cov}(\Delta \ln(c_{i,t}), \cdot)$	0.100***	-0.041***	-0.003	-0.002
	(0.005)	(0.003)	(0.002)	(0.002)
Observations	2,776	2,776	2,776	1,634

Notes: Non-durable consumption and net income, including transfers, are detrended by household characteristics using the same demographic and time controls employed in the estimation of the BPP pass-through coefficients presented in Figure 6. Standard errors clustered at the household level are shown in parenthesis. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

The first line in Table C.2.1 reports that the average autocovariance of log labor income growth is statistically significant up to lag $t + 2$, indicating that income growth is well approximated by a low-order moving-average process, close to MA(1), with only limited higher-order dependence (possibly reflecting time aggregation or measurement error) and limited persistence in transitory shocks. Consistent with this, their effects on consumption are short-lived. The second line reports higher-order cross-covariances between consumption growth and future income growth, while the third line reports autocovariances of consumption growth. In both cases, the estimates are not statistically significant beyond one period (two years), indicating that although consumption departs from a strict random walk, it does not exhibit persistent long-run dynamics.

C.2.1 Anticipations

As discussed in the main text, households may adjust current consumption in anticipation of future income innovations. Relaxing the random-walk restriction on consumption growth does not require assuming that households have full knowledge of the entire income process. However, advance information may be correlated with measured income risk—for instance, if part of the estimated risk reflects heterogeneity in information rather than in shocks. Table C.2.1 is informative about the horizon of anticipation: cross-covariances between consumption growth at t and future income growth are small and statistically insignificant beyond one period (two years), suggesting that any anticipatory effects are short-lived.

To account for one-period-ahead information, suppose that the permanent innovation at time t can be decomposed into two orthogonal components, η^s_{it} and η^a_{it}, where η^s_{it} is learned at time t and affects income at t, while η^a_{it} is learned at time $t-1$ but affects income only at t. The permanent component then evolves as

$$z_{it} = z_{i,t-1} + \eta^s_{it} + \eta^a_{it}, \qquad E(\eta^s_{it}) = E(\eta^a_{it}) = 0, \qquad \eta^s_{it} \perp \eta^a_{it}.$$

Analogously, the transitory innovation can be written as

$$\varepsilon_{it} = \varepsilon^s_{it} + \varepsilon^a_{it}, \qquad E(\varepsilon^s_{it}) = E(\varepsilon^a_{it}) = 0, \qquad \varepsilon^s_{it} \perp \varepsilon^a_{it},$$

where ε^s_{it} is unanticipated at $t-1$ and ε^a_{it} is known at $t-1$ but realized at t.

Throughout this appendix (as in the main text), $\psi^{\varepsilon}_{\text{BPP}}$ denotes the *BPP pass-through coefficient* for the transitory shock. With biennial data and the permanent–transitory decomposition in log income levels, income growth can be written as

$$\Delta y_{it} = \eta_{it} + (\varepsilon_{it} - \varepsilon_{i,t-1}),$$

so that Δy_{it} is mechanically MA(1). A key implication is that the lead autocovariance used by BPP is negative:

$$cov(\Delta y_{it}, \Delta y_{i,t+1}) = -var(\varepsilon_{it}) < 0.$$

Accordingly, the BPP pass-through can be written as the covariance ratio

$$\psi^{\varepsilon}_{\text{BPP}} = \frac{cov(\Delta c_{it}, \Delta y_{i,t+1})}{cov(\Delta y_{it}, \Delta y_{i,t+1})},$$

and the associated degree of partial insurance is $1 - \psi^{\varepsilon}_{\text{BPP}}$.

In the presence of one-period-ahead information about (part of) the permanent innovation, Kaplan and Violante (2010) show that the pass-through estimator may be biased because current consumption growth can covary with anticipated permanent news. In particular, one can write

$$\psi^{\varepsilon}_{\text{BPP}} = \psi_{\varepsilon} + \frac{cov(\Delta c_{it}, \eta^a_{i,t+1})}{cov(\Delta y_{it}, \Delta y_{i,t+1})}.$$

Using $cov(\Delta y_{it}, \Delta y_{i,t+1}) = -var(\varepsilon_{it})$, this becomes

$$\psi^{\varepsilon}_{\text{BPP}} = \psi_{\varepsilon} - \frac{cov(\Delta c_{it}, \eta^a_{i,t+1})}{var(\varepsilon_{it})}.$$

To obtain a simple approximation, suppose that a fraction α of the variance of permanent income innovations is known one period in advance, so that $var(\eta^a_{it}) = \alpha\, var(\eta_{it})$, and let ψ_{η^s} denote the pass-through of unanticipated permanent income innovations to

consumption. Approximating $cov(\Delta c_{it}, \eta^a_{i,t+1}) \approx \psi_{\eta^s}\, var(\eta^a_{it})$ yields

$$\psi^{\varepsilon}_{\text{BPP}} \approx \psi_{\varepsilon} - \alpha\, \psi_{\eta^s} \left[\frac{var(\eta_{it})}{var(\varepsilon_{it})}\right]. \tag{36}$$

Using expectations from the 1995 and 1998 waves of the SHIW, Kaufmann and Pistaferri (2009) estimate a pass-through of the permanent shock equal to 0.9341 once anticipation is accounted for, and report that about 30% of the variance of permanent earnings growth is known one period ahead. Table C.2.2 reports GMM estimates obtained by imposing these values in (36). For consistency with the required lags, the estimation sample is extended to include the 1991, 1993, and 1995 SHIW waves.[29] Under this configuration, the implied anticipated share of transitory income shocks is sizable (around one half). The pass-through for low-risk households remains higher and statistically significant, although precision is somewhat reduced, while the gap between high- and low-cash households widens. Results are robust to alternative discretionary values for α and for the imposed permanent-shock pass-through; these cases are available upon request.

C.2.2 Measurement Errors

With biennial data, transitory shocks realized early in the two-year window may raise consumption growth initially and partially reverse later within the same window. This time-aggregation mechanism attenuates the covariance between Δc_{it} and the instrumented income-growth variation used by BPP, implying that the BPP coefficient $\psi^{\varepsilon}_{\text{BPP}}$ should be interpreted as a *lower bound on pass-through* (and, equivalently, an *upper bound on the degree of insurance* $1 - \psi^{\varepsilon}_{\text{BPP}}$). More generally, exogenous borrowing constraints or interest-rate heterogeneity correlated with past shocks can also generate additional negative correlation between consumption growth and transitory-income innovations, reinforcing downward bias in pass-through. Accordingly, Table C.2.2 reports estimates that explicitly correct the moment conditions for measurement error. Using evidence for the SHIW from Kaufmann and Pistaferri (2009), I allow measurement error to account for roughly 30% of observed income variation.[30] The corrected estimates preserve the qualitative ranking: low-risk households display higher pass-through (lower insurance) than high-risk households, although with reduced precision, and the low- vs. high-cash gap widens. Results are robust to alternative assumptions, including cases in which the variance of measurement error is allowed to be as large as the variance of two-year transitory shocks. These cases are available upon request.

[29] The year 1991 is included for lag construction. In the 1991–2022 estimation sample, implied MPCs are computed as $\widehat{\text{MPC}}_g = \widehat{\psi}^{\varepsilon}_{\text{BPP},g} \cdot \widehat{\text{APC}}_g$, where $\widehat{\text{APC}}_g \equiv E[c/y \mid g]$ is recomputed on the same estimation subsample. This yields MPC point estimates of 0.31 for high-risk households and 0.46 for low-risk households. Likewise, the implied MPC is 0.34 for low-cash households and 0.43 for high-cash households. See Table C.2.4 for the corresponding pass-through coefficients.

[30] Specifically, Kaufmann and Pistaferri (2009) suggest that transitory income variation can be decomposed into an anticipated component (about 50%), an unanticipated component (about 20%), and measurement error (about 30%).

Table C.2.2: Anticipations and Measurement Errors

	Anticipation (50%)		Measurement Errors (30%)	
	low-risk	high-risk	low-risk	high-risk
ψ_j	0.773***	0.612***	1.080***	0.849***
	(0.068)	(0.071)	(0.096)	(0.098)
$\Delta\psi_j$	0.161*		0.231*	
	High-cash	Low-cash	High-cash	Low-cash
ψ_j	0.830***	0.545***	1.155***	0.543***
	(0.070)	(0.066)	(0.099)	(0.092)
$\Delta\psi_j$	0.285***		0.393***	

Notes: The table reports GMM estimates of the BPP pass-through coefficients. Households are classified as *low-risk* (*high-risk*) if their income risk measure lies below (above) the median of the CV distribution as presented in the main text. Similarly, *low* (*high*) resources denotes households below (above) the median of the cash-on-hand distribution. The row $\Delta\psi_j$ reports the difference in estimated coefficients between groups. The left columns account for advance information, while the right columns account for measurement error, both using estimated measures of bias from Kaufmann and Pistaferri (2009). For consistency with this paper, the sample includes observations from the 1993 and 1995 waves. Standard errors are robust to heteroskedasticity and clustered at the household level. Estimation allows for moment conditions to be computed on different subsamples, and does not impose a common estimation sample across moments. The weighting matrix is clustered by household and initialized as the identity matrix. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

C.2.3 BPP with various AR(1) persistent components

Following the BPP tradition, I restricted the income process to have either fully permanent or transitory shocks. However, shocks can be persistent, though not permanent. Accordingly, I consider the case in which the permanent component now is a persistent one, following an AR(1) with $\rho < 1$.

$$z_{it} = \rho z_{i,t-1} + \eta_{it}$$

Following Kaplan and Violante (2010), I now define the quasi-difference of log income as $\Delta\tilde{y}_t \equiv y_t - \rho y_{t-1}$. The identification of the two insurance coefficients can be achieved by setting $g_t^{\epsilon}(y) = \Delta\tilde{y}_{t+1}$ and $g_t^{\eta}(y) = \rho^2\Delta\tilde{y}_{t-1} + \rho\Delta\tilde{y}_t + \Delta\tilde{y}_{t+1}$. For the transitory shock, the estimator now is given by:[31]

$$\text{cov}(\Delta\tilde{y}_{it}, \Delta\tilde{y}_{i,t+1}) = -\rho\,\text{var}(\epsilon_{it}),$$

$$\text{cov}(\Delta c_{it}, \Delta\tilde{y}_{i,t+1}) = -\rho\,\text{cov}(\Delta c_{it}, \epsilon_{it}),$$

The first row of Table C.2.3 reports the baseline results underlying Figure 6 in the main text and compares them with results obtained under three alternative values of ρ 0.87,

[31] For the persistent shock, $\text{cov}(\Delta\tilde{y}_{it}, \rho^2\Delta\tilde{y}_{i,t-1} + \rho\Delta\tilde{y}_{i,t} + \Delta\tilde{y}_{i,t+1}) = \rho\,\text{var}(\eta_{it})$,
and $\text{cov}(\Delta c_{it}, \rho^2\Delta\tilde{y}_{i,t-1} + \rho\Delta\tilde{y}_{i,t} + \Delta\tilde{y}_{i,t+1}) = \rho\,\text{cov}(\Delta c_{it}, \eta_{it})$. See Kaplan and Violante (2010) for a detailed discussion on obtaining these estimators in the case of persistence.

Table C.2.3: Pass-through with Permanent and Persistent Income shocks

	Risk		Cash-on-Hand	
	high-risk	low-risk	Low-Cash	High-Cash
ψ_j, $\rho = 1$	0.352^{***}	0.643^{***}	0.453^{***}	0.587^{***}
	(0.069)	(0.056)	(0.064)	(0.079)
$\Delta\psi_j$		0.291^{***}		0.134
ψ_j, $\rho = 0.87$	0.379^{***}	0.700^{***}	0.531^{***}	0.568^{***}
	(0.091)	(0.085)	(0.100)	(0.134)
$\Delta\psi_j$		0.321^{**}		0.04
ψ_j, $\rho = 0.91$	0.382^{***}	0.697^{***}	0.527^{***}	0.572^{***}
	(0.088)	(0.084)	(0.095)	(0.129)
$\Delta\psi_j$		0.315^{**}		0.015
ψ_j, $\rho = 0.95$	0.384^{***}	0.694^{***}	0.522^{***}	0.576^{***}
	(0.086)	(0.082)	(0.095)	(0.125)
$\Delta\psi_j$		0.310^{**}		0.045

Notes: The table reports estimates of the pass-through coefficient to transitory shocks under the assumption that shocks are permanent (Figure 6) or persistent. Households are classified as *low-risk* (*high-risk*) if their income risk measure lies below (above) the median of the CV distribution as presented in the main text. Similarly, *low* (*high*) resources denotes households below (above) the median of the cash-on-hand distribution. The row $\Delta\psi_j$ reports the difference in estimated coefficients between groups. Statistical significance is based on a formal test of equality of the two group-specific estimates. In all cases, coefficients are estimated via an IV specification with group-specific first stages and time and individual fixed effects. Two-way standard errors clustered at the household and time levels are shown in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

0.91, and 0.95. Beyond these, I also verified that results are unaffected for a wide range of plausible values. These cases are available upon request. The results are consistent with the baseline estimates shown in Figure 6: The pass-through is significantly lower for high-risk households and indistinguishable between high- and low-cash households.

C.2.4 Robustness of BPP Pass-through to Discretionary Choices

Table C.2.4: Robustness tests on baseline pass-through coefficients

	Risk			Resources		
	High	Low	$\Delta\psi_j$	Low	High	$\Delta\psi_j$
Baseline (Figure 6)	0.353***	0.643***	0.289***	0.453***	0.588***	0.136
	(0.069)	(0.056)	(0.068)	(0.064)	(0.079)	(0.102)
No risk transitions	0.386***	0.720***	0.334***			
	(0.088)	(0.083)	(0.109)			
high-risk: temporary workers	0.202*	0.566***	0.364***			
	(0.114)	(0.056)	(0.130)			
Non-hand-to-mouth	0.394***	0.606***	0.212**	0.528***	0.513***	-0.015
	(0.071)	(0.067)	(0.085)	(0.125)	(0.057)	(0.142)
Income = total Income	0.360***	0.596***	0.236***	0.442***	0.559***	0.117
	(0.068)	(0.054)	(0.072)	(0.060)	(0.078)	(0.102)
Income = Wage Income	0.303***	0.458***	0.155**	0.346***	0.426***	0.080
	(0.068)	(0.076)	(0.068)	(0.051)	(0.122)	(0.129)
Total Consumption	0.361***	0.717***	0.355***	0.543***	0.583***	0.041
	(0.089)	(0.070)	(0.105)	(0.094)	(0.074)	(0.127)
Period: 1998-2016	0.404***	0.677***	0.273***	0.421***	0.654***	0.233**
	(0.074)	(0.073)	(0.099)	(0.064)	(0.083)	(0.101)
Period: 1991-2022	0.404***	0.656***	0.252***	0.429***	0.618***	0.189**
	(0.063)	(0.060)	(0.084)	(0.059)	(0.065)	(0.089)
Age: 18-72	0.383***	0.665***	0.282***	0.408***	0.660***	0.252***
	(0.063)	(0.070)	(0.083)	(0.064)	(0.070)	(0.085)
Period: 1989-2022, Age: 18-72	0.405***	0.650***	0.245***	0.411***	0.656***	0.246***
	(0.058)	(0.057)	(0.071)	(0.059)	(0.055)	(0.071)
Dropping household FE	0.443***	0.606***	0.163**	0.451***	0.607***	0.156
	(0.078)	(0.050)	(0.065)	(0.078)	(0.067)	(0.089)

Notes: Each row reports pass-through estimates under alternative specifications or sample definitions. Households are classified as *low-risk* (*high-risk*) if their income risk lies below (above) the median of the CV distribution; *Low* and *High* resources refer to the median of the cash-on-hand distribution. $\Delta\psi_j$ reports the difference between group-specific coefficients; significance is based on equality tests. The baseline reproduces Figure 6. Alternative rows modify risk transitions, define high risk using temporary contracts, restrict to non-hand-to-mouth households, redefine income (total or wage), use total consumption, alter the sample period, expand the age range, or drop household fixed effects. Estimates use IV with group-specific first stages, time and households fixed effects (except the last row), and two-way clustered standard errors (household and time). Significance: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

C.3 Details on 2020-2022 Sample

Table C.3.1: Descriptive Statistics SHIW Sample 2020-2022.

	Full sample		2020		2022	
	Mean	SD	Mean	SD	Mean	SD
Demographics						
Male	0.62	0.49	0.60	0.49	0.63	0.48
Age	47.02	6.41	47.90	6.09	46.49	6.54
Married	0.65	0.48	0.64	0.48	0.66	0.48
Years of education	1.81	0.40	1.75	0.43	1.84	0.37
Resident in the South	12.15	20.74	11.33	20.24	12.65	21.04
Size of the family	3.05	1.28	3.03	1.32	3.07	1.25
Number of labor income earners	1.70	0.65	1.62	0.64	1.75	0.65
City size $\geq$ 500,000	0.15	0.36	0.06	0.24	0.20	0.40
Financials (×1,000)						
Disposable labor income	48.88	25.91	42.72	21.93	52.59	27.39
Net Cash-on-hand	56.73	165.97	43.70	168.47	64.58	164.08
Net liquid wealth	52.54	165.09	39.83	167.75	60.20	163.13
Net illiquid wealth	243.67	818.00	234.90	1221.63	248.96	417.47
Non-durable consumption	31.19	19.96	27.17	13.73	33.62	22.59
Average Propensity to Consume	0.70	0.30	0.70	0.28	0.69	0.30
Average Self-Reported MPCs	0.45	0.30	0.44	0.29	0.45	0.30
Income Risk						
Coefficient of Variation (CV)	0.04	0.03	0.05	0.03	0.04	0.03
Number of observations	957		360		597	

Notes: The table reports descriptive statistics for the SHIW sample over 2020–2022. The "Full sample" column pools the two waves, while the remaining columns report wave-specific means and standard deviations. Financial variables are expressed in thousands of euros. Net liquid wealth includes deposits and other liquid financial assets net of short-term liabilities. Net illiquid wealth includes real estate and other non-liquid assets net of debt. Following Jappelli and Pistaferri (2014, 2020), net cash-on-hand is defined as net liquid wealth plus one month of disposable income (i.e., annual disposable income divided by 12). The coefficient of variation (CV) measures household income risk as described in the main text.

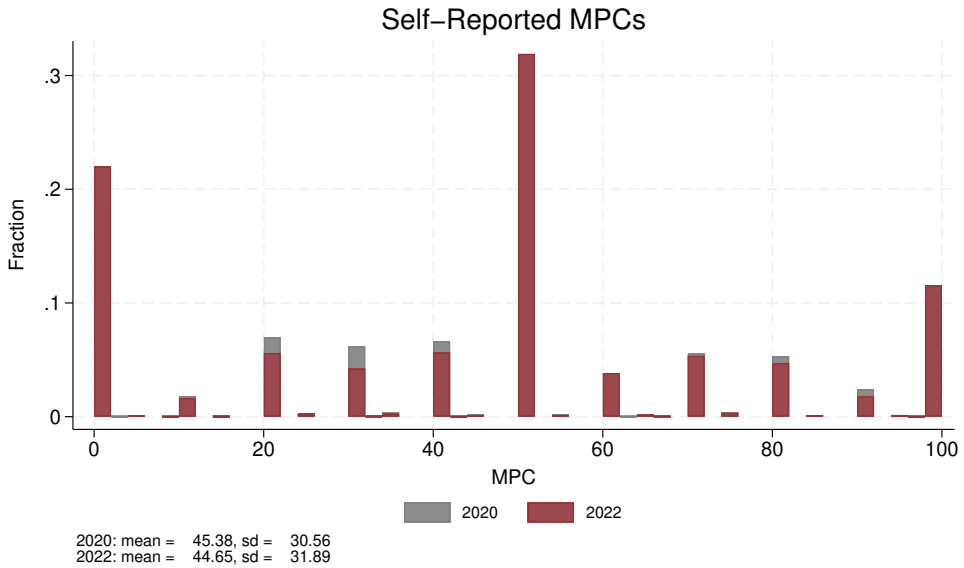


Figure C.3.1: Self Reported MPCs: 2020-2022

Notes: The panel reports the distribution of Self-reported MPCs for the two waves, 2020 and 2022.

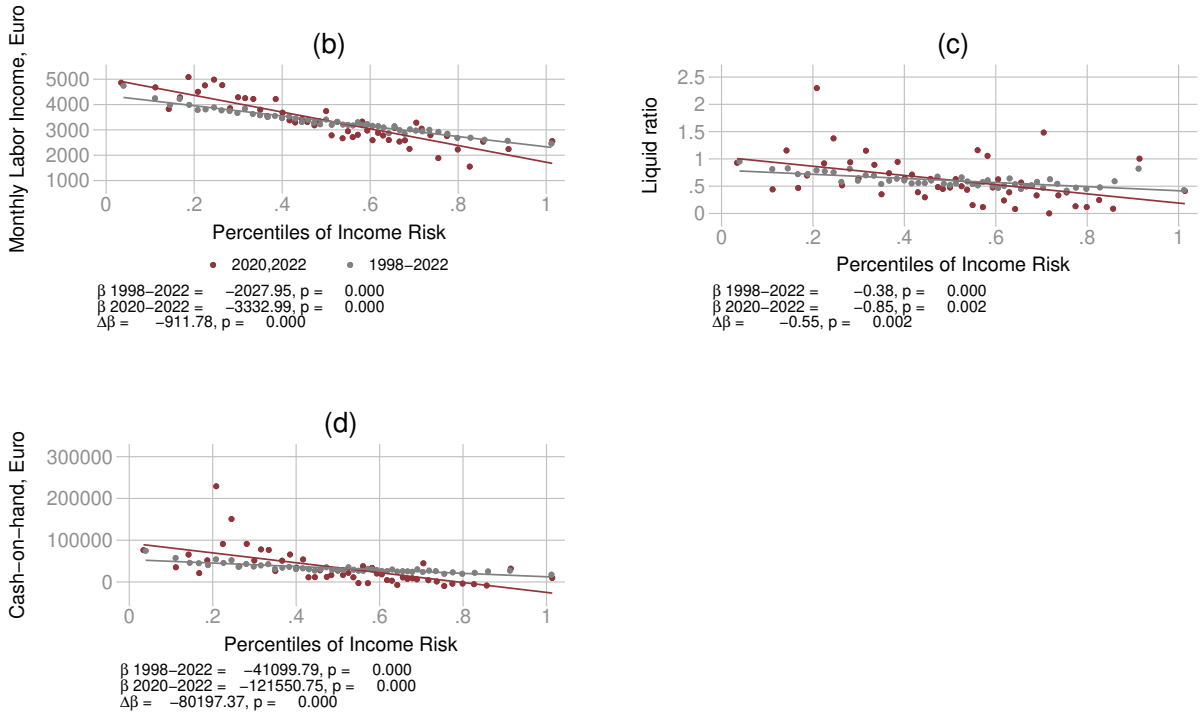


Figure C.3.2: Self Reported MPC and income risk covariates across samples

Notes: The top panels overlap the correlates of income risk reported in Figure 3 with the resulting relationship for the sample 2020-2022. Regressions control for age, sex, household size, education, marital status, region, and year fixed effects. The reported coefficient corresponds to income risk; statistical significance and p-values are reported in each panel's note.

D Details on income risk

D.1 Empirical implementation of Step 1

This section documents the empirical implementation of the income risk measure defined in the main text. The object of interest is the individual-level coefficient of variation of unpredictable income,

$$CV_{i,t} \equiv \frac{\mathbb{E}(|Y_{i,t} - \mathbb{E}(Y_{i,t} \mid X_{i,t})| \mid X_{i,t})}{\mathbb{E}(Y_{i,t} \mid X_{i,t})},$$

where $Y_{i,t}$ denotes disposable labor income and $X_{i,t}$ represents the household information set.

First stage: conditional mean. The conditional mean of income $\mathbb{E}(Y_{i,t} \mid X_{i,t})$ is estimated using a Poisson pseudo–maximum likelihood model with exponential mean specification,

$$\mathbb{E}(Y_{i,t} \mid X_{i,t}) = \exp(X_{i,t}'\beta).$$

Predicted values $\widehat{\mathbb{E}}(Y_{i,t} \mid X_{i,t})$ are censored at the maximum observed income in the sample to avoid extrapolation in the upper tail of the income distribution.

Table D.1.1: First-stage estimation: conditional mean of income

	First stage
Estimator	Poisson PML (exponential mean)
Dependent variable	$Y_{i,t}$
Number of observations	9,779
Controls $X_{i,t}$	Micro and macro covariates
Age and age^2 interactions	Yes
Regional fixed effects	Yes
Robust standard errors	Yes
Log pseudo-likelihood	−985,951.1
Pseudo-R^2	0.979

Second stage: conditional dispersion. Given the estimated conditional mean, I construct the absolute deviation

$$D_{i,t} = \left| Y_{i,t} - \widehat{\mathbb{E}}(Y_{i,t} \mid X_{i,t}) \right|.$$

The conditional expectation of $D_{i,t}$ is estimated as

$$\mathbb{E}(D_{i,t} \mid X_{i,t}) = \exp(X_{i,t}'\gamma),$$

using the same information set $X_{i,t}$ adopted in the first stage.

Table D.1.2: Second-stage estimation: conditional mean absolute deviation

	Second stage
Estimator	Poisson PML (exponential mean)
Dependent variable	$\lvert Y_{i,t} - \widehat{\mathbb{E}}(Y_{i,t} \mid X_{i,t})\rvert$
Number of observations	9,779
Controls $X_{i,t}$	Same as first stage
Age and age^2 interactions	Yes
Regional fixed effects	Yes
Robust standard errors	Yes
Log pseudo-likelihood	−6,876,221.9
Pseudo-R^2	0.308

The fitted values from the two stages are combined to obtain the empirical counterpart of $CV_{i,t}$ defined above.

Table D.1.3 summarizes the set of covariates included in $X_{i,t}$. Controls are grouped into micro-level determinants—capturing income and labor-market history as well as household demographics—and macro-level conditions at the national and regional level. In addition to the baseline controls reported in the table, the empirical specification allows for flexible heterogeneity by age through interactions of each control with age and age squared, as implemented in the estimation code.

Table D.1.3: CV Predictors

Micro		**Macro**
Income	**Demographics**	**National, Regional**
Log lagged income $(t-1)$	Age	Unemployment rate (U)
Log lagged income squared	Age squared	GDP growth
Log lagged income cubed	Education: secondary or higher	$U\times$ South regions
Indicator: lagged income $= 0$	Male	GDP growth $\times$ South regions
Months worked $(t-1)$	Married	Region fixed effects (ISTAT)
Indicator: positive income while unemployed	Couple with children	
Full-year employment indicators $(t-1, t-2, t-3)$	Large city (>500,000 inhabitants)	
Extra hours worked $(t, t-1)$	South regions	
Lagged unemployment income $(t-1)$		
Permanent contract $(t-1)$		
No part-time employment $(t-1)$		

Notes: The table reports baseline controls included in $X_{i,t}$. In the estimation, each micro- and macro-level control is also interacted with age and age squared. Region fixed effects are defined using ISTAT region codes and are likewise interacted with age and age squared.

D.2 Other Correlates of Income Risk

Figure D.2.1 presents a set of correlations between income risk and key household characteristics. Panel A shows that the coefficient of variation (CV) is positively correlated with its lagged value, indicating persistence in income risk. Panel B documents that households working the full year exhibit substantially lower income risk than those with fewer working days. Panel C shows a negative relationship between income risk and lagged income levels. As discussed in the main text, these patterns are consistent with the findings in Arellano et al. (2022).

Panel D further shows that income risk is negatively correlated with permanent income, constructed as average income over the households' observed lifespan. Panel E reports a negative association between income risk and age. Finally, Panel F provides little support for the hypothesis that households facing higher income risk adjust along the intensive margin by supplying more hours, as the relationship between hours worked and income risk is essentially flat.

All panels in Figure D.2.1 report fractional polynomial regressions of income risk (measured by the CV) on the variables shown on the horizontal axis. Grey bands denote 95% confidence intervals. The results are robust to estimating a linear specification saturated with controls, individual and regional fixed effects, and time fixed effects.

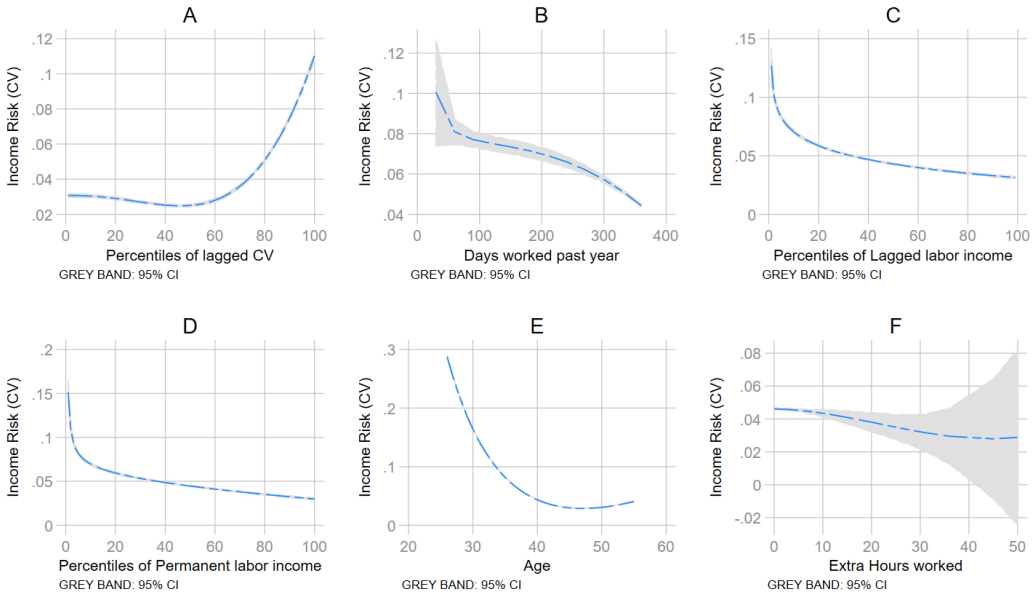


Figure D.2.1: Other Correlates of Income risk

D.3 Income risk decomposition

Figure D.3.1 reports the decomposition of labor-income variance for the two groups. A key limitation of the coefficient of variation (CV) used in the main text is that it captures overall income dispersion without distinguishing between permanent and transitory sources of risk. In this section, I address this limitation by separately estimating the permanent and transitory components of income risk for low-risk and high-risk households,

as defined in the main text. This decomposition allows me to assess whether broadly defined income risk reflects differences in both long-run and short-run income variability.

Framework. I estimate a covariance-structure model for residual log earnings. After removing observable characteristics, log earnings for individual i in cohort c at time t are written as the sum of a permanent and a transitory component:

$$y_{i,c,t} = u_{i,t} + \varepsilon_{i,t}. \tag{37}$$

The permanent component follows a random walk:

$$u_{i,t} = u_{i,t-1} + w_{i,t}, \tag{38}$$

where $w_{i,t}$ is an innovation with variance σ_w^2 and $E(u_{i,t-1}w_{i,t}) = 0$. This specification implies that permanent shocks accumulate over time and generate persistent income dispersion. The transitory component is allowed to display short-run serial correlation and is modeled as an AR(1) process:

$$\varepsilon_{i,t} = \rho_\varepsilon \varepsilon_{i,t-1} + \epsilon_{i,t}, \tag{39}$$

where $\epsilon_{i,t}$ is an innovation with variance σ_ϵ^2 and $E(\epsilon_{i,t}\epsilon_{i,s}) = 0$ for $t \neq s$. This specification corresponds to the Restricted Income Profiles (RIP) benchmark: individuals differ in realized permanent and transitory shocks but do not exhibit heterogeneous deterministic life-cycle growth rates (i.e., no random growth component). In particular, the data do not support the Heterogeneous Income Profiles (HIP) extension with individual-specific growth rates correlated with initial earnings.[32]

Connection with the model in Section 5. The specification above is an auxiliary empirical decomposition used to match the autocovariance structure of income growth in the data, separately by type. In the quantitative model of Section 5, income is instead approximated by a parsimonious process with a persistent AR(1) component and an i.i.d. transitory shock. More precisely, I choose $(\sigma_{\eta,j}, \sigma_{\varepsilon,j})$ in Section 5 so that, at the SHIW biennial frequency and for each type j, the model matches the two one-step second moments of residual income growth implied by the auxiliary RW+AR(1) decomposition above, namely $\mathrm{Var}(\Delta y_{i,t})$ and $\mathrm{Cov}(\Delta y_{i,t}, \Delta y_{i,t+1})$; higher-lag autocovariances are close to zero in the data and therefore not targeted. Accordingly, the serial correlation in the transitory component estimated in (39) is not imposed literally in the structural model; rather, its implications for the one-lag covariance of $\Delta y_{i,t}$ are absorbed by the persistent AR(1) component of the approximation, while the remaining short-run variation is captured by an i.i.d. transitory shock. This mapping reconciles the auxiliary RW+AR(1)

[32]The model is estimated following Doris et al. (2011), which matches the empirical covariance matrix of income residuals to its theoretical counterpart implied by the random-walk permanent component and AR(1) transitory component. A Wald test does not reject the constant permanent-factor restriction, supporting the RIP benchmark. Allowing for an ARMA(1,1) specification of the transitory component does not materially improve the fit in this biennial setting, so I adopt the parsimonious random-walk-plus-AR(1) specification.

decomposition with the AR(1)+i.i.d. income process used for simulation in Section 5 without re-estimating the empirical moments.

Results. Figure D.3.1 shows that high-risk households (above the median CV) exhibit higher variance in both the permanent and the transitory components of income. Hence, the broader income-risk measure used in the main text reflects differences in both long-run income uncertainty and short-run income fluctuations.

Importantly, the variances reported in Figure D.3.1 refer to the dispersion of realized income innovations within each risk group, conditional on observable characteristics. They do not represent the variance of expected income growth, but rather the variance of the stochastic components driving income dynamics.

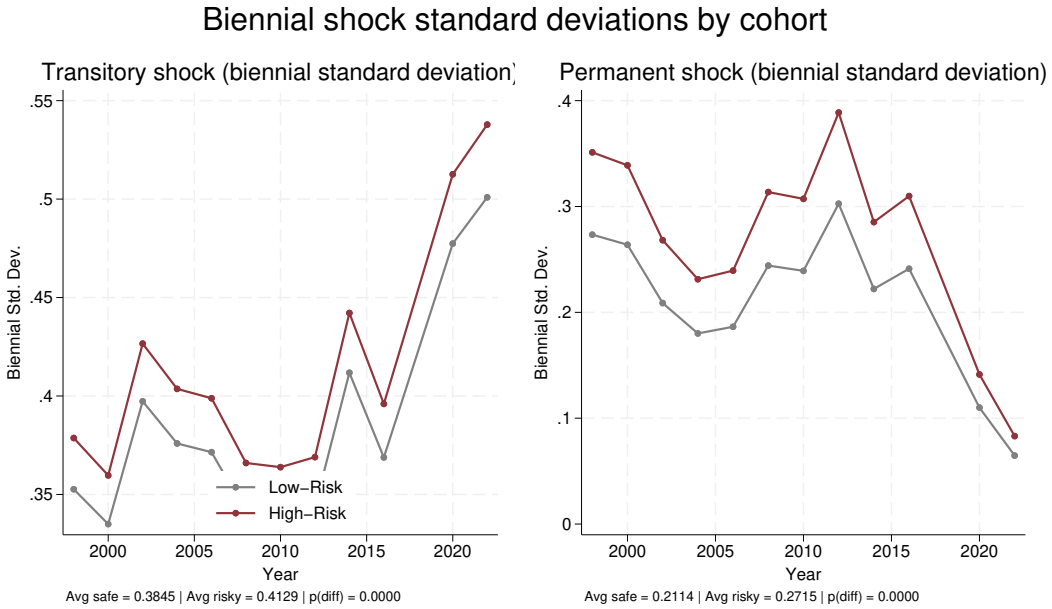


Figure D.3.1: Biennial Variances of Permanent and Transitory Income Shocks by Risk Group