

Heterogeneity in MPC Beyond Liquidity Constraints: The Role of Permanent Earnings

Jeanne Commault*

June 2024

Abstract

The marginal propensity to consume (MPC) is a central object in economics that is key to understand the transmission of shocks. Recent empirical findings challenge the standard view that its distribution is mostly explained by constraints on liquid wealth: (i) some people with substantial liquid wealth have a high MPC; (ii) higher current earnings, which should relax the constraints, do not reduce the MPC. I note that, in the standard consumption model, it is the combination of people's liquid wealth and the variance of their future earnings that determines their precautionary motive and constraints, thus their MPC. Everything else being equal, a higher permanent component of earnings, means a higher variance of future earnings and a higher MPC. This can explain (i)-(ii). Survey data support a large and positive effect of permanent earnings on the MPC. Numerical simulations can replicate those findings quantitatively in a model with realistic earnings risk.

Key words: Marginal Propensity to Consume, Earnings Risk, Precautionary Saving, Heterogeneous Agent Model

JEL: D11, D12, D15, E21

1 Introduction

People's marginal propensity to consume out of a one-time unexpected income shock (MPC) is a central object in economics. Its magnitude determines the extent to which aggregate shocks are

*Department of Economics, Sciences Po, 28 Rue des Saints-Pères, 75007 Paris, France; jeanne.commault@sciencespo.fr. I am grateful to participants at the Sciences Po Lunch seminar, Princeton Macro Lunch, NYU Macro Lunch, Federal Reserve Bank of Chicago, Conference on Individual Risk and the Economy, Minnesota Macro Conference for helpful comments. I specifically thank Jean-Marc Robin, Fabien Tripier, and Gianluca Violante for helpful comments and suggestions.

passed onto demand and can transmit to the rest of the economy. Through the same mechanism, it determines the effectiveness of fiscal and monetary policies.

The current narrative about why, on average, people respond to a one-time unexpected shock centers around liquid wealth. Most of those who have little liquid wealth are constrained thus respond strongly to a shock. The others smooth out a shock over many periods and do not respond strongly. Because empirically quite a lot of people keep little liquid wealth and hold most of their wealth in illiquid form, this view can explain why the average MPC is substantially above zero,

While useful and consistent with a number of observations, the influence of liquid wealth on the MPC is not sufficient to explain two empirical facts that the more recent literature documents: (i) some people with medium or high levels of liquid wealth have non-zero MPCs, so the extent to which liquid wealth reduces the MPC remains empirically modest; and (ii) higher current earnings do not associate with lower MPCs, despite their positive impact on people's immediately available cash, which should relax liquidity constraints; higher current earnings can even associate with higher MPCs everything else being equal. These two facts cannot be explained by the particular behavior of durable consumption because they also hold for MPCs associated with nondurable spending.

In this paper, I note that another potentially important determinant of the MPC is future earnings risk. Now, earnings typically include a permanent component, that is, a factor to which all current and future realizations of earnings are proportional. An increase in this permanent component raises the variance of future earnings and the magnitude of the earnings shocks that people face. Comparing consumers with the same liquid wealth but a different level of permanent earnings, those with a higher permanent earnings face larger future shocks but hold the same accumulated wealth. As a result, they have a stronger precautionary motive and a higher MPC. This positive effect of the permanent component of earnings on the MPC can explain the two stylized facts above. I find empirical support for this theoretical prediction: at a given level of liquid wealth, a higher level of permanent earnings correlates with a higher MPC in US survey data. The effect of permanent earnings is of a magnitude comparable to that of wealth. Numerically, a standard incomplete market model can generate the same positive relation between permanent earnings and the MPC as in survey data and generate the two stylized fact that motivate the analysis, provided the earnings risk is realistic enough. A more realistic specification of earnings also raises the fraction of people with a high MPC.

My first contribution is analytical. I consider a standard macroeconomic consumption model to illustrate the mechanism at the core of the paper. There is no exogenous borrowing constraint, only a natural borrowing constraint that never binds, to show that the mechanism holds even absent

binding constraints. The earnings process includes a permanent component, to which all current and future earnings realizations are proportional. A higher level of this permanent component increases total expected resources, but it multiplies up the consumers' risky human capital without multiplying up the risk-free liquid wealth they have accumulated. As a result, people with a higher level of permanent earnings save a larger share of their (higher) earnings, to accumulate more liquid wealth. A windfall income gain that directly provides liquidity relaxes this need for saving more for them. Conversely, they cut their consumption more in response to an unexpected one-time loss because such a drop in liquidity raises their need for saving more than it does to others. Their MPC is higher.

This positive relation between permanent earnings and the MPC can explain the two stylized facts above. The people who have more liquid wealth are more likely to have a large permanent component of earnings, which partly offsets the negative effect of liquid wealth on their MPC. They can still have a relatively high MPC, explaining fact (i). Also, an increase in current earnings may not reduce the MPC if this increase is partly driven by the permanent component of earnings, explaining fact (ii).

My second contribution is empirical: I establish that, in US survey data, conditional on wealth and demographics, a higher level of permanent earnings correlates with a higher MPC. To measure this, I first design a method to recover an empirical counterpart to permanent earnings. The typical difficulty is that the underlying permanent component of earnings is not directly observed: surveys only report total earnings. I use expected future earnings to identify this permanent component: expectations have been used to decompose income shocks, I show that they can also be used to identify the level of permanent earnings of the respondents. I implement this method in the New York Fed Survey of Consumer Expectations (SCE). In line with the earnings specification, the variance of future earnings increases with my measure of permanent earnings.

In a reduced form regression analysis, permanent earnings correlate significantly and positively with people's reported MPC out of a hypothetical one-time change in resources. Quantitatively, a one standard-deviation increase in permanent earnings correlates with a 0.04 level increase in the yearly MPC for total consumption, including spending on both durable and nondurable goods. I observe two measures of MPC, one based on a question about a negative one-time income shock and one based on a question about a positive one-time income shock. The result is robust to the choice of the MPC: permanent earnings correlates positively with both. Regarding the ability of this correlation to explain the stylized facts (i) and (ii), I find that the impact of permanent earnings is almost as large as that of liquid wealth, thus large enough to partly offset it. It is also large enough so that total current earnings can have a only a negligible effect on the MPC.

One policy implication is that, in order to maximize the consumption response, targeting stimulus checks to people based on their income is not the most effective. This is because income includes components that both raise and reduce the MPC. The ideal would be to be able to observe permanent income and wealth. Absent this possibility, a targeting based on both age and income can raise the average MPC of the targeted group more than when a targeting based on income only.

My third contribution is to show that numerical simulations of a standard model can replicate both my survey results and the two stylized facts that motivate the analysis. I find that modeling earnings as proposed in Guvenen, Karahan, Ozkan, and Song (2021) is key. More precisely, when I shift from their process to a simple transitory-permanent process with normally distributed shocks, as is typical in numerical simulations, the effect of permanent earnings on the MPC no longer offsets enough the effects of liquid wealth and of the transitory component of earnings. Such simulations are then unable to generate the stylized facts (i) and (ii).

Additionally, besides making the model able to generate stylized facts (i) and (ii), incorporating this richer structure of earnings risk raises the average MPC out of a positive shock. This is important because the empirical observation of a large average MPC out of a positive shock is a crucial stylized fact that motivated the shift away from representative agent models of consumption, unable to generate such a MPC, and opened the exploration of MPC heterogeneity. Similar to the distinction between liquid and illiquid wealth, the earnings risk that people face is an observable dimension that can be calibrated externally and transparently introduced in models. The combination of the two dimensions raises the MPC out of a positive shock close to the high level that people report, itself consistent with MPC estimates from natural experiments.

Related literature. My overall results are most closely related to the empirical literature that examines the characteristics that covary with people's MPCs. This is the literature that established the stylized facts that constitute my starting point. What the present paper brings to this literature is an explanation for these stylized facts, and an additional stylized fact that supports this explanation, which is that in survey data the permanent component of earnings has a significant, positive, and large effect on the MPC. More precisely, the stylized fact (i) that motivate my analysis and that became an important finding in this literature is that, besides demographics, the main characteristic that consistently affects the MPC is liquid wealth. The extent to which it reduces the MPC, however, is relatively modest (see e.g. Jappelli and Pistaferri (2014), Baker (2018), Aydin (2019), Ganong, Jones, Noel, Farrell, Greig, and Wheat (2020), or Fagereng, Holm, and Natvik (2021)).¹

¹On this, Fuster, Kaplan, and Zafar (2020) write '... the only observable characteristic that has been robustly shown to correlate with MPCs is holdings of liquid wealth, and even then the explanatory power of wealth for MPC heterogeneity is weak.', p1.

In contrast, the stylized fact (ii) is that the effect of total current earnings is typically not significant (see e.g. Parker, Souleles, Johnson, and McClelland (2013), Broda and Parker (2014), Misra and Surico (2014), Fagereng, Holm, and Natvik (2021), Parker, Schild, Erhard, and Johnson (2022), Boutros (2022)). Some papers even find a positive and significant effect of current earnings on the MPC. That is the case of Kueng (2018)² and of Lewis, Melcangi, and Pilossoph (2022). In Jappelli and Pistaferri (2014), the authors document a significant and negative correlation between current income and the MPC out of a shock proportional to people's income, which is not incompatible with a positive correlation between current income and the MPC out of a shock of the same size for all.

My theoretical result builds on and extends the scope of the few studies that have examined the role of the permanent component of earnings in the consumer's problem. Among them, Carroll (2006) notes that one can normalize the consumer's problem by the permanent component of earnings to solve it with one less variable. This is useful in particular when simulating such a model. The implication is that what matters for the ratio of consumption to permanent earnings is the value of the ratio of wealth to permanent earnings. I expand the insight to show that the normalization does not just yield a new problem with one less variable, but in fact keeps the problem the same because the consumption function is homogeneous. I note that this has implications for the MPC, that is, the response to a one-time shock. Relatedly, Carroll (2009) considers the marginal propensity to consume out of a permanent shock (MPCP), and shows that it is smaller than one in a standard consumption model. Straub (2019) examines the effect of permanent earnings on this MPCP, that is, the concavity of consumption in permanent earnings.

Finally, my last result is in the same spirit as Kaplan and Violante (2014). Their study shows that modeling wealth more realistically, by distinguishing between liquid and illiquid wealth, makes a fraction of people with high wealth more responsive to a windfall income change because their high wealth is mostly illiquid. I show that modeling earnings risk more realistically, by including a richer specification, makes a fraction of people with high liquid wealth more responsive to a windfall income change because they are exposed to high earnings risk.

²The paper of Kueng (2018) examines the response to an anticipated income gain (not an unexpected shock), and proposes a mechanism that is specific to anticipated changes. The mechanism I identify can explain why people with higher earnings respond more to an unexpected shock, and it can bolster the mechanism proposed by Kueng (2018) as to why people with higher earnings respond more to an anticipated shock upon realization and not upon learning about it.

2 Permanent earnings and the MPC in a standard model

2.1 An income-fluctuation model with a transitory-permanent process

Model To present the intuition of how an increase in the permanent component of earnings raises a consumer's MPC in a standard framework, I consider a simple income-fluctuation problem with a transitory-permanent earnings process and only one asset.

A consumer i is finite-lived, with T the length of their life. The consumer chooses consumption expenditures at period t , denoted c_t^i , to maximize lifetime expected utility subject to a number of constraints

$$V_t^i(a_t^i, e^{p_t^i}, e^{\varepsilon_t^i}) = \max_c u(c) + \beta E_t \left[V_{t+1}^i(a_{t+1}^i, e^{p_{t+1}^i}, e^{\varepsilon_{t+1}^i}) \right] \quad (2.1)$$

$$\text{with Utility conditions: } u'(\cdot) > 0, u''(\cdot) < 0, \text{ and } u'''(\cdot) > 0 \quad (2.2)$$

$$\text{Positive spending: } c > 0, \quad (2.3)$$

$$\text{Budget constraint: } a_{t+1}^i = (1+r)a_t^i + y_t^i - c, \quad (2.4)$$

$$\text{Earnings: } y_t^i = e^{p_t^i} e^{\varepsilon_t^i} \quad (2.5)$$

$$\text{Permanent component: } e^{p_{t+1}^i} = e^{p_t^i} e^{\eta_{t+1}^i}, \quad (2.6)$$

$$\text{Terminal wealth: } a_{T+1}^i \geq 0. \quad (2.7)$$

Utility is time-separable and at each period depends only on contemporaneous consumption. The period utility function $u(\cdot)$ is such that marginal utility is positive, decreasing, and convex in consumption: $u'(\cdot) > 0$, $u''(\cdot) < 0$, and $u'''(\cdot) > 0$. This implies that people are prudent, so uncertainty pushes them to save more than they would have otherwise. The marginal utility $u'(c)$ approaches infinity when consumption c approaches zero. The discount factor β captures how much consumers discount utility between two consecutive periods.

The positive spending condition (2.3) imposes that consumption be strictly positive at each period.

The budget constraint (2.4) states that, to store their wealth from one period to the next the consumer only has access to one risk-free liquid asset. The term a_t^i denotes the level of this asset at the beginning of period t —or at the end of $t-1$. The risk-free return rate is r . This rate r is such that $\beta(1+r) \leq 1$.

The labor earnings specification, described with (2.5) and (2.6), is a transitory-permanent process: earnings are the product of a permanent component $e^{p_t^i}$ that evolves as a multiplicative random walk and of a transitory innovation $e^{\varepsilon_t^i}$ that is an i.i.d. shock. Because the permanent

component e^{p_i} multiplies the value of the permanent component at the next period, it multiplies each realization of earnings until the rest of the consumer's lifetime: at $t + s$, earnings are $y_{t+s}^i = e^{p_i} e^{\eta_{t+1}^i + \dots + \eta_{t+s}^i} e^{\varepsilon_t^i}$. It thus plays the role of a scaling factor. Note that this specification encompasses an even simpler specification in which the permanent component is just a multiplicative fixed effect $e^{p_i} = e^{p^i}$. This is for instance the specification in Straub (2019). Incidentally, the transitory-permanent process has initially been used to model the earnings of individuals (e.g. in Meghir and Pistaferri (2004)) but is now used more broadly to model the net income of households, including the effect of taxes and transfers (e.g. in Blundell, Pistaferri, and Preston (2008) or in numerical simulations). In this theoretical part, I assume for simplicity that earnings and net income coincide—there are no taxes nor transfers. In the empirical and numerical part, the transitory-permanent process models earnings. For the precautionary motive to be strictly positive, I impose that people face a strictly positive amount of transitory earnings uncertainty: $var(\varepsilon) > 0$.

The terminal condition on wealth (2.7) states that the consumer cannot die with a strictly positive level of debt: assets at the end of the last period T —and the beginning of $T + 1$ —have to be non-negative. The combination of this condition with the period budget constraints and positive spending constraints generates a natural borrowing constraint that prevents people from holding a level of debt superior to what they could ever repay. This constraint never binds because marginal utility approaches infinity as consumption approaches zero: consumers would never put themselves in the situation of possibly consuming zero in the future. In the remainder of the section, I drop the household index i to ease notations.

MPC definition What I want to examine is the partial effect on consumption of an unexpected one-time income shock. Such a shock would be modeled as an unexpected term w_t entering the budget constraint at t such that $a_{t+1} = (1 + r)a_t - c_t + y_t + w_t$. This equation shows that the term w_t would have the same effect on the consumption decision as an unexpected change in the beginning-of-period wealth term a_t : in this model an unexpected shock w_t such as a lottery win has the same impact on consumption as an unexpected change in a_t such as an unexpected inheritance. As a result, the MPC is equivalently defined as the partial effect of a_t on c_t :

$$MPC_t \equiv \frac{\partial c_t}{\partial a_t}.$$

2.2 The effect of the permanent component on the MPC

The main result that I prove is that, in this framework, for a class of utility functions that encompasses the standard isoelastic case, the permanent component e^{p_t} has a strictly positive effect on

the MPC

$$\frac{\partial MPC_t}{\partial e^{p_t}} > 0.$$

This is a straightforward but overlooked implication of two results that are already around in the literature, at least in some special cases: that consumption is concave in wealth, and that consumption is homogeneous of degree one in wealth and permanent earnings. For the proof, I define precautionary saving PS_t as the difference between what the consumers save and what they would save under perfect foresight at t , if they solved exactly the same problem as described by (2.1)-(2.7) except that, from period t on, income was equal to its expected value at t with probability one and the product of the discount factor and interest rate $R = \beta(1+r)$ was equal to one.

Theorems. In the model described above by (2.1)-(2.7), at any period $t < T - 1$

- (1) When the ratios of temperance over prudence and of prudence over risk-aversion are both non-increasing, consumption is concave in wealth. This means that the MPC is lower at a higher level of asset a_t

$$\frac{\partial MPC_t}{\partial a_t} = \frac{\partial^2 c_t}{\partial a_t^2} < 0$$

When absolute prudence and absolute risk-aversion are both non-increasing (and one of them strictly), then the reason why consumption is concave in wealth is because precautionary saving decreases with wealth but less so at a higher level of wealth. This means that the MPC is higher than it would under perfect foresight, and the gap between the two is smaller at a higher level of wealth

$$\frac{\partial PS_t}{\partial a_t} < 0 \text{ and } \frac{\partial^2 PS_t}{\partial a_t^2} > 0 \text{ thus } MPC_t = \frac{\partial c_t^{PF_t}}{\partial a_t} - \frac{\partial PS_t}{\partial a_t} > \frac{\partial c_t^{PF_t}}{\partial a_t} \text{ and } \frac{\partial MPC_t}{\partial a_t} = -\frac{\partial^2 PS_t}{\partial (a_t)^2} < 0.$$

- (2) When utility displays constant relative risk-aversion, consumption is homogeneous of degree one in risk-free liquid wealth a_t and permanent earnings e^{p_t} so by Euler's theorem it is equal to the weighted sum of its derivatives with respect to a_t and e^{p_t} . This means that the MPC is homogeneous of degree zero in a_t and e^{p_t}

$$c_t = a_t \frac{\partial c_t}{\partial a_t} + e^{p_t} \frac{\partial c_t}{\partial e^{p_t}} \text{ thus } 0 = a_t \frac{\partial MPC_t}{\partial a_t} + e^{p_t} \frac{\partial MPC_t}{\partial e^{p_t}}.$$

In addition, when utility only displays non-increasing relative risk-aversion and relative prudence, I still have $c_t \leq a_t \frac{\partial c_t}{\partial a_t} + e^{p_t} \frac{\partial c_t}{\partial e^{p_t}}$.

- (3) When utility displays constant relative risk-aversion, the MPC is higher at a higher level of permanent earnings e^{p_t} for consumers with strictly positive risk-free liquid wealth $a_t > 0$

$$\frac{\partial MPC_t}{\partial e^{p_t}} = -\frac{a_t}{(e^{p_t})} \frac{\partial MPC_t}{\partial a_t} > 0 \text{ when } a_t > 0.$$

In addition, when utility only displays non-increasing relative risk-aversion and relative prudence and the ratios of temperance over prudence and prudence over risk-aversion are both non-increasing, then around low levels of permanent earnings, I also have $\frac{\partial MPC_t}{\partial e^{p_t}} > 0$.

Proof and intuition for Theorem 3. The ideas that consumption is concave in wealth, and homogeneous of degree one in wealth and permanent earnings are present in the literature. My proofs of Theorem 1 and 2 extend existing results to a broader range of utility functions (Theorem 1) and a more general definition of permanent earnings (Theorem 2). My main contribution lies in establishing, with Theorem 3, that those familiar results have a direct implication for the effect of permanent earnings on the MPC. I present the intuition for Theorems 1 and 2 and show how Theorem 3 stems from them. The proofs of Theorems 1 and 2 are in Appendix A.

The intuition for Theorem 1 is that, under perfect foresight, people consume a fixed fraction of their total expected resources. A windfall income gain raises consumption only because it raises total expected resources: following such a gain, people consume the same fixed fraction of a larger amount. In the presence of uncertainty, when absolute prudence and risk aversion are positive and non-increasing, the MPC is higher than under perfect foresight because, on top of increasing total resources, a windfall income gain provides liquidity that relaxes the need to save and accumulate liquidity for precautionary reasons. This decrease in the need for saving generates a higher increase in consumption. Now, when the ratios of temperance over prudence and prudence over risk-aversion are non-increasing, the magnitude by which precautionary saving decreases with a windfall income gain is smaller at a higher level of wealth. As a result, the MPC is smaller at a higher level of wealth.

The intuition for Theorem 2 is that there are two types of resources that can be used to finance consumption in the model: the risk-free wealth that consumers have accumulated, and the flow of current and future earnings they expect, which are subject to shocks and which are scaled by the level of permanent earnings. When wealth and permanent earnings are multiplied by k , the two types of resources are multiplied by k and the ratio of risk-free to risky resources remains the same. When utility displays constant relative risk aversion, people keep saving and consuming the

same fraction of these resources when their ratio of risk-free to risky resources remains the same. Optimal consumption is thus multiplied by k when both types of resources are. Consumption is homogeneous of degree one in wealth and permanent earnings. The MPC measures the effect of an increase in resources that tilts the ratio of risk-free to risky resources up. A one unit increase in liquid wealth tilts the ratio by $1/k$ times less when wealth and permanent earnings are both k times larger. By homogeneity, consumption over permanent earnings increases by $1/k$ times less. Thus, consumption, in level, increases by the same amount. The MPC is homogeneous of degree zero: it remains the same when wealth and permanent earnings are multiplied by the same constant.

Combining those results, I obtain Theorem 3. Because the MPC depends on the ratio of wealth to permanent earnings, a higher level of permanent earnings has the same impact on the MPC as a lower level of wealth when wealth is strictly positive: $\frac{\partial MPC_t}{\partial e^{p_t}} = -\frac{a_t}{(e^{p_t})} \frac{\partial MPC_t}{\partial a_t}$. Because a lower level of wealth raises the optimal level of precautionary saving and makes it more sensitive to one-time shocks, $\frac{\partial MPC_t}{\partial a_t} < 0$, consumption is also more sensitive to one-time income shocks at a higher level of permanent earnings: $\frac{\partial MPC_t}{\partial e^{p_t}} > 0$.

To keep things simple, Theorem 3 presents the case with perfect homogeneity. However, this perfect homogeneity is not always necessary. In the limit where permanent earnings e^{p_t} approaches zero, current and future earnings approach zero. People consume only out of their accumulated assets, there is no earnings thus no uncertainty, and the consumer's problem approaches its perfect foresight counterpart. In that case, consumption is homogeneous of degree one in wealth and permanent earnings and equal to the weighted sum of its derivatives. In the second part of Theorem 2, I prove that, for a broad range of utility functions, at any level of permanent earnings (always strictly positive), consumption is smaller than the weighted sum of its derivatives : $c_t \leq a_t \frac{\partial c_t}{\partial a_t} + e^{p_t} \frac{\partial c_t}{\partial e^{p_t}}$. This means that, consumption increases less with an increase in permanent earnings away from zero than the weighted sum of its derivatives does

$$a_t \frac{\partial^2 c_t}{\partial a_t \partial e^{p_t}} + \frac{\partial c_t}{\partial e^{p_t}} + e^{p_t} \frac{\partial^2 c_t}{\partial (e^{p_t})^2} > \frac{\partial c_t}{\partial e^{p_t}} \quad (2.8)$$

$$\frac{\partial^2 c_t}{\partial a_t \partial e^{p_t}} > -\frac{e^{p_t}}{a_t} \frac{\partial^2 c_t}{\partial (e^{p_t})^2} > 0 \quad (2.9)$$

I prove in Commault (2024) that consumption is strictly concave in permanent earnings under the same conditions required for concavity in wealth. Therefore, $(\partial^2 c_t)/(\partial (e^{p_t})^2) < 0$ and it has to be that the MPC increases with permanent earnings for the weighted sum of derivatives to increase more than consumption with permanent earnings.

Discussion. Regarding Theorem 1, my proof extends the result of Carroll and Kimball (1996) who show, without conditions on $\beta(1+r)$, that the concavity of consumption in wealth holds when utility displays HARA. The HARA condition means that the ratio of prudence over risk-aversion and the ratio of temperance over prudence are constant—together with all higher-order ratios. I show here that, under the realistic restriction that $\beta(1+r) \leq 1$, it is sufficient to let these two ratios be non-increasing. The result is consistent with the findings of Toda (2021). His paper shows that, in the standard consumption model with an increasing and concave utility function, for the consumption function to be concave regardless of the value of the β , $(1+r)$, and regardless of the distribution of the strictly positive income, the utility function must display HARA. If the utility function of the model additionally has to display positive prudence ($u'''(\cdot) > 0$), as I impose here, the utility function must display CRRA, not just HARA, for the consumption function to be concave. The paper of Gong, Zhong, and Zou (2012) shows that this HARA requirement can be relaxed in a deterministic consumption problem (without any uncertainty) with $\beta(1+r) \geq 1$ at each period. What I show here is that the CRRA assumption can be largely relaxed, beyond HARA, under the realistic restriction that $\beta(1+r) \leq 1$.

The fact that precautionary saving decreases with wealth is a straightforward implication of the concavity of consumption in wealth when $\beta(1+r) = 1$. This is because, in this case, consumption is linear in wealth when removing earnings uncertainty, and consumption approaches asymptotically from below its no-uncertainty counterpart as wealth approaches infinity. When consumption is concave in wealth, the precautionary gap between its value and its no-uncertainty counterpart value must widen as wealth decreases away from its asymptotic limit: precautionary saving decreases when wealth decreases. Now, in the general case with $\beta(1+r) \leq 1$, apart from the HARA utility case, removing earnings uncertainty is no longer sufficient to make consumption linear in wealth. For instance, without earnings uncertainty, consumption is concave when the ratio of prudence over risk-aversion is strictly decreasing and $\beta(1+r) < 1$. That is why I take as a benchmark a situation without earnings uncertainty and where $\beta(1+r) = 1$. I prove that the gap between consumption in general and in this benchmark is decreasing in wealth but this is no longer a straightforward implication of concavity since the general case and the benchmark case no longer necessarily converge when wealth approaches infinity.

Regarding Theorem 2, the idea that consumption-related variables can be expressed as functions of the ratio of wealth to permanent earnings exists in the literature, mostly in the context of numerical simulations, to simplify the model and speed up its computation. More precisely, Carroll (2006) and Carroll (2009) develop the pioneer insight, that, in an income-fluctuation model with a CRRA utility it is possible to divide consumption, wealth, and total earnings by current perma-

permanent earnings (referred to as a normalization by permanent earnings) and obtain a new consumers' problem with one less state variable. My result generalizes this insight by proving homogeneity. This means proving that the policy function is the same in the initial problem as in the normalized problem: if the solution of the main problem is $c_t = f(a_t, e^{p_t}, e^{\varepsilon_t})$ then $c_t/e^{p_t} = f(a_t/e^{p_t}, 1, e^{\varepsilon_t})$. This is different from proving there exists a \tilde{f} such that the solution c_t of the problem verifies $c_t/e^{p_t} = \tilde{f}(a_t/e^{p_t}, 1, e^{\varepsilon_t})$. The difference is not important for numerical simulations but homogeneity is useful for building analytical results on the behavior of consumption in those models. Coming from another strand of the literature, Straub (2019) establishes homogeneity, not just the possibility to normalize, in a similar model as me but with a more restrictive earnings process and assumptions about initial wealth than I do. His paper shows that consumption is homogeneous of degree one in a_t and e^{p_t} when the permanent component of earnings is a time-invariant multiplicative fixed effect that is not subject to any shock and when initial wealth is zero.³ I show that one does not need the permanent component to be fixed nor initial risk-free wealth to be proportional to it for consumption to be homogeneous in risk-free wealth and permanent earnings. Importantly, these papers did not consider the implication of this possibility to normalize or of homogeneity for the relation between permanent earnings and the MPC.

Regarding Theorem 3, Wang, Wang, and Yang (2016) derive equilibrium conditions on the consumption of solution of a continuous-time and infinitely lived income-fluctuations model with Epstein-Zin preferences. Simulating numerically the model, they suggest that introducing shocks to permanent earnings can raise the average MPC.

Extension to exogenous borrowing constraints. This theoretical section relies on a simple model to illustrate the mechanism I identify and prove Theorem 3. However, I can extend the model on some dimensions. The mechanism I exhibit, and the result of Theorem 3, can extend to a situation with an exogenous borrowing constraint under some conditions on the life-cycle trend of earnings. In the presence of an exogenous constraint, there are two regimes: one in which the constraint binds, and one in which it does not. When the constraint binds in the same way before and after a one-time shock, two consumers that are similar except that one has a higher level of permanent earnings have the same MPC: because they are constrained they consume all of a one-time gain or have to reduce their consumption by all of a one-time loss. When the constraint does not

³Note that Straub (2019)'s Proposition 1 is a proportionality result. It is similar to Theorem 3, thus not inconsistent with Theorem 4. It states that, if one observes an individual with twice as high permanent earnings, then their consumption will be twice as high. This is because, from the additional assumptions of time-invariant permanent earnings and initial wealth proportional to permanent earnings, current wealth is always proportional to permanent earnings, thus twice as large for an individual with permanent earnings twice as large. Current consumption is thus also twice as large.

bind before and after a one-time shock, Theorem 3 applies and a consumer with a higher level of permanent earnings has a higher MPC. Thus, Theorem 3 holds, not strictly, for consumers away from the threshold.

It would also hold, not strictly, if the wealth threshold above which the exogenous constraint no longer binds did not change with a consumer's permanent earnings. Yet, a higher level of permanent earnings changes the region where an exogenous constraint binds. In the absence of any life-cycle trend in income, that is, when the expected value of income is the same at all periods, an increase in permanent earnings reduces the wealth threshold above which people are no longer constrained. That is because, denoting c_t^* the consumption that would be chosen absent constraints and \bar{L}_{t+1} the exogenous borrowing limit on wealth at the beginning of $t + 1$, people are constrained when $a_{t+1} = (1 + r)a_t + y_t - c_t^* < \bar{L}_{t+1}$. Because optimal consumption does not increase as much as current earnings with an increase in permanent earnings in the absence of life-cycle trend (see Carroll (2009)), I have $\frac{\partial c_t^*}{\partial e^{pt}} < \frac{\partial y_t}{\partial e^{pt}}$. Thus, a marginal increase in permanent earnings reduces the gap between constrained consumption and optimal consumption. In that case, Theorem 3 might not hold at some levels of wealth: a consumer who is constrained can have a higher MPC than a similar consumer with a higher permanent earnings who is not constrained thanks to its higher permanent earnings. However, when there is an upward-sloping trend to earnings, so that the expected value of future earnings is higher than the expected value of current earnings, it is possible to have $\frac{\partial c_t^*}{\partial e^{pt}} > \frac{\partial y_t}{\partial e^{pt}}$. In that case, a consumer with higher permanent earnings remains constrained over a wider range of wealth values than a consumer with lower permanent earnings. There can then be two reasons why consumers with a higher permanent earnings respond more to a one-time income shock: because the same exogenous constraint is as or more binding for them or because they have a stronger precautionary motive.

Regarding borrowing constraints, other papers have looked, not at the effect of permanent earnings on the strength of the constraint, but at the related questions of the effect of liquidity constraints, their tightening, and the addition of more constraints on the shape of the consumption function, in the absence of risk (Holm (2018)) and in the presence of risk (Carroll, Holm, and Kimball (2021)), for HARA utility functions.

Extension to capital income. The current model does not include capital income and as such is not adequate to model the top of the income and wealth distribution. Note however that the result would extend to a framework in which, instead of doing some salaried work, people can be entrepreneurs and the profits from entrepreneurial work have an individual-specific, permanent, component. The heterogeneous returns literature documents the existence of such an individual-

specific component to capital income (see e.g. Bach, Calvet, and Sodini (2020), Fagereng, Guiso, Malacrino, and Pistaferri (2020) and Smith, Zidar, and Zwick (2022)). Then, if investment is somehow constrained and cannot be adjusted easily, the same mechanism that I exhibit here applies: at the same level of wealth, somebody with higher individual-specific component to capital income is more exposed to capital income risk and has a stronger desire to accumulate safe assets. A windfall gain stimulates consumption more because it partly fills this desire for safe assets and reduces the need to achieve it via saving.

3 Measuring permanent earnings in survey data

3.1 The Survey of Consumer Expectations

Survey. To test empirically this theoretical prediction, I use data from the Survey of Consumer Expectations (SCE) of the Federal Reserve Bank of New-York. It is a monthly online survey with a rotating panel of about 1,300 household heads based in the US. A household head is defined as a person in the household who owns, is buying, or rents the home. A household may have multiple co-household heads. Respondents stay on the panel for up to twelve months before rotating out of the panel. The survey started in June 2013. While the Core Survey takes place monthly, its topical modules only take place either every four months or every year. Because they are reported in different modules, I only observe income, consumption and wealth around the same period once every year, at the end of the year. This means that there is no panel dimension in the analyses that include earnings, consumption and wealth. There is, however, a panel dimension in some analyses of earnings. I describe the way in which I match these different modules in B.1. Also, because not all the modules started in 2013, the period over which I observe jointly these variables is from the end of 2015 to the end of 2018.⁴

Earnings. I obtain current annual earnings, expected future annual earnings, and the probability to be employed at the next period from questions in the Labor Market module of the SCE. From this module, I also observe the probability that respondents assign to the occurrence of earnings-changing events in the future, such as receiving job offers of different amounts or becoming self-employed. This makes it possible to build a measure of the variance of future earnings as foreseen by the individuals themselves.

⁴See Armantier, Topa, Klaauw, and Zafar (2017) for technical background information on the SCE, and www.newyorkfed.org/microeconomics/sce.html for additional information.

MPCs. I build the MPCs out of negative and positive transitory shocks from questions in the Household Spending module of the SCE. The survey asks respondents to consider a hypothetical situation in which their annual household income next year would be 10% higher: 'Suppose next year you were to find your household with 10% more income than you currently expect. What would you do with the extra income?'. They are also asked about a hypothetical situation in which their annual household income next year would be 10% lower: 'Now imagine that next year you were to find yourself with 10% less household income. What would you do?'. The survey elicits the response in two steps, which helps avoid most responses being at 0, 0.5, or 1. In a first step, respondents are offered the choice between corner solutions—using it all to spend, saving it all, or using it all to repay debt—or stating they would combine between these three. A majority of respondents declare they would combine. In a second step, those stating they would combine are asked to quantify what percentage of the shock would be absorbed by each of the three channels. Respondents receive this question in December, so 'next year' corresponds to immediate future. Because the questions are about shocks that are in percentage of household income, I control for household income in the analyses, to control for the size of the shocks people are asked about.

In the survey, the fact that the shock is a one-time event, meaning that income decreases or increases next year but not afterwards, is suggested but not strongly stressed. For this reason, I check whether the answers are consistent with responses from other surveys in which the hypothetical shock is explicitly described as a one-time occurrence. I find that the responses are similar.⁵ This suggests that households do seem to interpret the question in the SCE in the same way they interpret a question about an explicitly one-time income shock. In their study, Koşar, Melcangi, Pilossoph, and Wiczer (2023) also consider the response to the question that I use as measuring the response to a one-time shock. Furthermore, they compare it to the respondents' reported use of the one-time COVID stimulus check (which they have access to). They find the response broadly similar to the response to the hypothetical question that I use.

Now, a fact that emerges from the consumption literature is that, after initially using part of a one-time income gain to repay some of their debt, people take on new debts and go back to a level

⁵In Fuster, Kaplan, and Zafar (2020) the question is 'Now consider a hypothetical situation where you unexpectedly receive a one-time payment of [\$500,\$2,500,\$5,000] today. We would like to know whether this extra income would cause you to change your spending behavior in any way over the next 3 months.', which includes the word 'one-time'. The options are then similar to those in the SCE: first people report whether what they would do, and then they are asked about the exact percentages. The average share of \$2,500 and \$5,000 extra income that would be used for spending (excluding debt repayment) over the next three months are 0.11 and 0.14 (Table A6). In my sample, the average share of a 10% annual income gain that would be used for spending (excluding debt repayment) over the next year is 0.15, thus very close. In Crossley, Fisher, Levell, and Low (2021), the wording is almost the same as in Fuster, Kaplan, and Zafar (2020). The share of a \$500 gain that would be used for spending over the next three months is 0.11 (Table 1).

of debt close to what it was before the income gain. Agarwal, Liu, and Souleles (2007) find that debt decreases over the two months following a windfall income gain, and then increases again, so that nine months after the shock, people no longer have less debt than before the shock. The point estimate even implies they have more debt nine months after a positive shock than they would have otherwise but it is not significantly different from zero. This means that, nine months after the shock, most of the initial debt repayment is transformed into extra consumption financed by new debt that would not have been taken up if the previous debt had not been repaid. Since what I look at is the MPC over the following year, not over the next two to three months, this can be an issue. To overcome it, I categorize a certain share of what people report as debt repayment as an increase in consumption. I set this share to match the marginal propensity to repay debt out of a positive shock over the following year of 7% estimated in Fagereng, Holm, and Natvik (2021), where debt is directly observed. This yearly time horizon is the same as the one over which people are asked to consider the increase in income in the SCE. Note that, because very few people answer that they would take on new debt to cope with a negative shock, this adjustment has virtually no impact on the MPC out of a negative shock. The analyses with the MPC out of a negative shock are thus robust to making this adjustment or not.

The resulting MPC out of a positive shock that I obtain is 0.462, well within range of MPCs obtained in natural experiments. The recent findings of Orchard, Ramey, and Wieland (2023) and Borusyak, Jaravel, and Spiess (2022), which correct for potential biases, put the MPC out of a positive shock around 0.30 for total consumption and 0.10 for strictly nondurable consumption. A reported MPC out of a positive shock of 0.462 over a year for total consumption is consistent with an MPC of 0.30 over a quarter. Fagereng, Holm, and Natvik (2021) estimate the yearly MPC of total consumption out of a lottery win to be on average 0.523. The MPC out of a negative shock that I obtain is 0.788. The latter is more difficult to compare to natural experiment values because natural experiments of an income loss are more scarce.

Regarding the validity of the MPC out of hypothetical shocks versus the MPC out of actual shocks, Parker and Souleles (2019) compare the two and find them to yield the same average MPC and to comove. Also, as mentioned when discussing the one-time nature of the shock, Koşar, Melcangi, Pilossoph, and Wiczler (2023) also observe both the hypothetical and reported MPCs of the SCE respondents. They obtain similar results with both measures.

Wealth. I build wealth from a question in the Housing module of the SCE asking respondents to select which category of net non-housing wealth their household belongs to among fourteen possible bins, ranging from below five hundred dollars to above one million dollars.

Demographics. I use demographic characteristics in two instances: to net out their effects from expected future annual earnings and build permanent earnings, and to control for their effect on the MPC. The earnings-related demographic variables that I use are dummies for the respondents' gender, age group, educational attainment, willingness to take risks, and year dummies. The consumption-related demographic variables that I use to detrend consumption are the same, excluding gender because I move from the individual to the household level, and adding six dummies for the number of household members (from one to more than six), one dummy for whether there is at least one child (below 18) in the household, eleven dummies for ranges of total household income, five dummies for labor force status of the spouse, one dummy for whether the household has a budget or plan for monthly spending), and dummies for US state of residence. I obtain all those variables from the Core module of the SCE.

Selection and CPI deflating. I exclude non-employed respondents from the sample. This does not mean that I am assuming away the risk of non-employment since I let those employed face future non-employment risk. The reason for excluding the non-employed is that, to build permanent earnings, I assume that people draw their earnings shocks from the same distributions conditional on demographics. This does not seem to hold when I include non-employed respondents, who appear to draw earnings shocks from riskier distributions. I further drop respondents with yearly earnings below \$1,885, following Guvenen, Karahan, Ozkan, and Song (2021). To abstain from modeling the education and retirement decision, I also select out people below age 25 and above age 55. Finally, I trim the top and bottom 1% of the expected future earnings, earnings, consumption, and variance of future earnings variables—re-coding the top and bottom values as non-reported so that the order in which I trim the variables does not matter. I drop the 22 people whose reported responses to the MPC questions (what they would do with the loss or gain) do not add up to 100% so I do not further trim the MPCs. I deflate all the \$ variables using the non-seasonally adjusted Consumer Price Index (CPI), and express them all in 2014\$. More precisely, I use as reference the average price index over the second half of the 2014 year. I present descriptive statistics of my main variables in this final sample in Appendix B.3.

3.2 Measuring permanent earnings

A general earnings process. I present a method to identify the permanent component of earnings that is consistent with a general model of earnings encompassing the simple transitory-permanent process that I use in section 2 as a special case. I let the annual earnings y_t^i of individual i at year t

follow the specification proposed in Guvenen, Karahan, Ozkan, and Song (2021). The authors find that such a specification fits well the moments of administrative US data on earnings⁶

$$\text{Annual earnings: } y_t^i = \underbrace{(1 - v_t^i)}_{\text{Employment status}} \underbrace{e^{\tilde{\alpha}^i + \zeta z^i}}_{\text{Fixed effect}} \underbrace{e^{p_t^i}}_{\text{Highly persistent}} \underbrace{e^{\varepsilon_t^i}}_{\text{Transitory}} \underbrace{e^{g(t)}}_{\text{Age trend}} \quad (3.1)$$

$$\text{Persistent component: } e^{p_t^i} = (e^{p_{t-1}^i})^\rho e^{\eta_t^i}, \quad (3.2)$$

$$\text{Nonemployment: } v_t^i \sim \begin{cases} 0 \text{ (employment) with prob. } 1 - p_{v_{t-1}^i}, \\ 1 \text{ (nonempl.) with prob. } p_{v_{t-1}^i}. \end{cases} \quad (3.3)$$

This expression states that annual earnings are the product of a dummy for employment status v_t^i , a fixed effect $e^{\alpha^i} = e^{\tilde{\alpha}^i + \zeta z^i}$ that includes a part depending on observed, non-time varying, demographics z^i , a persistent component $e^{p_t^i}$, a transitory innovation $e^{\varepsilon_t^i}$, and a deterministic age trend $e^{g(t)}$. The log of the persistent component evolves as an AR(1) process with η_t^i its innovation. In practice, this component is virtually permanent because Guvenen, Karahan, Ozkan, and Song (2021) estimate its persistence to be $\rho = 0.991$ and the length of the working age period is less forty years, not infinite. Because I do not seek to estimate this process but simply to find an empirical counterpart to the permanent component, I do not need to assume any specific distributions for these innovations. I can also let people with different demographic characteristics draw their innovations from different distributions. Specifically, I can let the mean and variance of the distributions of η_t^i and ε_t^i depend linearly on demographic and year dummies.

The non-employment dummy at t , v_t^i , is a one/zero dummy for whether the individual is mostly employed or mostly non-employed over the year.⁷ I directly observe the probability to be non-employed at t , $p_{v_{t-1}^i}$, from a survey question. I can therefore be agnostic about its specification.⁸

Defining permanent earnings. My objective is to capture the part of this earnings process that is

⁶Guvenen, Karahan, Ozkan, and Song (2021) identify two specifications that fit the data well, which are numbered (5) and (6) in their paper. Here, I draw from their specification (5).

⁷When Guvenen, Karahan, Ozkan, and Song (2021) let the non-employment dummy correspond to a length that they let estimate, they find it to be very close to one, which corresponds to one year. Although this value is likely to capture an average, it means that this assumption of 'mostly non-employed over the year' versus 'mostly employed' still generates a good fit between the specification and the moments of the administrative data.

⁸Because people respond at a given period about their probability to be non-employed at the next period, however, I assume that this probability of non-employment is entirely determined at $t - 1$. This differ slightly from Guvenen, Karahan, Ozkan, and Song (2021), where the probability to be non-employed at t is determined after the realization of the contemporaneous persistent earnings innovation η_t^i . Yet, I note that because I let the distribution of the persistent earnings shocks $e^{\eta_t^i}$ depend on individual characteristics at $t - 1$ and because I let $p_{v_{t-1}^i}$ also depend on individual characteristics at $t - 1$, my specification is still consistent with a correlation between $e^{\eta_t^i}$ and $p_{v_{t-1}^i}$.

akin to a scaling factor multiplying the realizations of current and future earnings. In this specification, such a factor corresponds to the product of the highly persistent component $e^{p_t^i}$ (virtually permanent) and of the fixed effect at the average demographics value $e^{\tilde{\alpha}^i + \zeta \bar{z}}$. I additionally normalize the value of this product. Indeed, the effect of a one unit change in $e^{\tilde{\alpha}^i + \zeta \bar{z}} e^{p_t^i}$ is not directly interpretable in terms of earnings change. I re-scale it so that a one one unit change in my definition of permanent earnings corresponds to a one dollar increase in annual earnings, at the average age trend and average realization of the current transitory innovation. I denote $perm_t^i$ the re-scaled permanent component

$$perm_t^i = e^{\tilde{\alpha}^i + \zeta \bar{z}} e^{p_t^i} e^{\bar{\epsilon}} e^{\bar{g}} \quad (3.4)$$

The bar over the variables denotes their average value in the sample.

Using detrended expected future earnings to measure permanent earnings. Under this general specification, dividing expected future annual earnings by the probability to still be employed at the next period and taking the log of the resulting term yields

$$\underbrace{\ln\left(\frac{E_t^i[y_{t+1}^i]}{(1-p_{vt}^i)}\right)}_{\text{Observed}} = \underbrace{\overbrace{\rho p_t^i + \tilde{\alpha}^i}^{res_t^i}}_{\approx 1} + \underbrace{\zeta z^i + \ln(E_t^i[e^{\eta_{t+1}^i}]) + \ln(E_t^i[e^{\epsilon_{t+1}^i}]) + g(t+1)}_{\text{Captured through demographic and year dummies}}. \quad (3.5)$$

If demographic and year dummies affect linearly the values of the mean and variance of the distributions from which people draw their transitory and persistent innovations, differences in $\ln(E_t^i[e^{\eta_{t+1}^i}])$ and $\ln(E_t^i[e^{\epsilon_{t+1}^i}])$ are captured by a linear regression over such dummies.⁹ Also, ρ is estimated by Guvenen, Karahan, Ozkan, and Song (2021) to be close to one. As a result, the residual res_t^i from a regression of $\ln(E_t^i[y_{t+1}^i]/(1-p_{vt}^i))$ on demographic dummies (for gender, age group, educational attainment, and willingness to take risks of the head) and year dummies coincides with $p_t^i + \tilde{\alpha}^i$.¹⁰ I re-scale and re-arrange this residual to obtain my measure of permanent earnings. I multiply this residual by the average log-income among employed respondents,

⁹This is because the log of the expected values $\ln(E_t^i[e^{\eta_{t+1}^i}])$ and $\ln(E_t^i[e^{\epsilon_{t+1}^i}])$ approximate as $\ln(E_t^i[e^{\eta_{t+1}^i}]) \approx E_t^i[\eta_{t+1}^i] + \text{Var}_t^i(\eta_{t+1}^i)/2$ and $\ln(E_t^i[e^{\epsilon_{t+1}^i}]) \approx E_t^i[\epsilon_{t+1}^i] + \text{Var}_t^i(\epsilon_{t+1}^i)/2$.

¹⁰It coincides with $p_t^i - \bar{p} + \tilde{\alpha}^i - \bar{\alpha}$ but I set the average sample value of $\bar{p} + \bar{\alpha}$ to zero without loss of generality (any non-zero constant can be captured in $e^{\zeta \bar{z}}$).

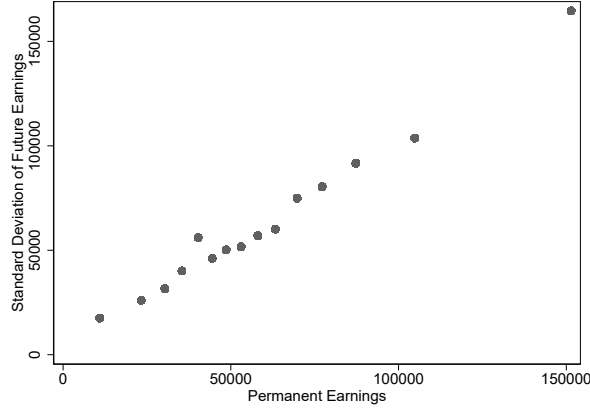


Figure 1: The Evolution of the Standard Deviation of Future Earning With Permanent Earnings

denoted $\overline{\ln(y)}|_{v_i=0} = \bar{\epsilon} + \zeta \bar{z} + \bar{g}$. I then take the exponential

$$e^{res_i^j \times \overline{\ln(y)}|_{v_i=0}} = e^{p_i^j + \bar{\alpha}^i} e^{\bar{\epsilon}} e^{\delta \bar{z}} e^{\bar{g}} = perm_t^i.$$

Table 5 in Appendix B.3 presents summary statistics on the raw variables that I use to build permanent earnings and on my resulting measure of permanent earnings.

The method relates to the papers of Pistaferri (2001) and Attanasio, Kovacs, and Molnar (2020), which use expectations to identify separately the transitory and permanent components of the shocks that people face. Here I use the same insight to identify the level of permanent earnings. Furthermore, I show that it can work for a general earnings specification when the probability of non-employment and demographics are observed. Other methods to proxy the level of permanent earnings include using simply current earnings or an average of past and current earnings. A method which, like mine, attempts to actually remove the transitory component from current earnings and allow for a general earnings specification with non-employment shocks, is that of Braxton, Herkenhoff, Rothbaum, and Schmidt (2021) who develop a filtering algorithm that can be used if a long enough time series is available.

Permanent earnings and the standard deviation of future earnings. As a first check of my theoretical mechanism, I examine how the standard deviation of future earnings varies with permanent earnings. One difficulty with measuring risk is that the ex-ante risk that an individual faces is not observed, only the one outcome that realizes ex-post is. To overcome this, I rely on questions in the SCE asking respondents about their estimated probabilities to receive job offers of different amounts, to still be working for the same employer, to still be employed but working for a different

employer, to be self-employed, and to be non-employed in the future. I use those to compute, for each respondent, the future earnings levels they consider possible and the probabilities they assign to each level. This then gives me an individual measure of the variance and standard deviation of future earnings.¹¹ I discretize people’s permanent earnings in thirty bins of equal size. I compute the average permanent earnings and the average standard deviation of future earnings of the respondents in each bin. Figure 1 plots the relation between the two. It shows that people’s standard deviation of future earnings increases linearly with their permanent earnings. This is consistent with an earnings process where permanent earnings multiplies the realizations of future earnings, thus multiplies their standard deviation.¹² The standard deviation is large because I consider earnings, not income, and earnings become zero in non-employment. Note that the questions I use to build the standard deviation of future earnings are not the expected future annual earnings question that I use to build permanent earnings so the linear relation I obtain is not built in.

I also check that the coefficient of variation of future earnings does not vary with current earnings, as predicted by my specification. I find that the coefficient of variation decreases with earnings when I include both non-employed and employed respondents, as in Arellano, Bonhomme, Vera, Hospido, and Wei (2021), but becomes flat when I exclude the non-employed—as in their paper as well where the decrease is driven by people with low attachment to the labor market. I present those results this in Appendix B.5.

Ruling out anticipations, using expected annual earnings at a short horizon, and other potential concerns One potential concern of using expectations to separate out the transitory component is that, if the realization of the future transitory component is expected, future expected earnings would include it. In that case, what I measure would not be $perm_t^i$ but $perm_t^i \times e^{\varepsilon_{t+1} - \bar{\varepsilon}}$. This is because people’s expectations would include the realized idiosyncratic value of $e^{\varepsilon_{t+1}}$, and it would not wash out in the predicted value of a regression over demographic and period dummies. I can however test for this. I look at the covariance between my measure of permanent earnings at t and the realized innovation to log-earnings at $t + 1$. If the transitory shock was anticipated, it would be present in both my measure of permanent earnings and the realized value of the shocks, and the two would covary. Contrary to that, I find that their correlation is not significant and small. I

¹¹ See Appendix B.2, which details the survey questions and the method that I use to build this variance.

¹² Incidentally, the specification (5.1) only implies a proportional and positive relation between permanent earnings and the standard deviation of future earnings *conditional on being employed in the future*. Figure 1 shows that the relation is there even between permanent earnings and the *unconditional* standard deviation of future earnings, it means that, although higher permanent earnings might protect from unemployment thus reduce future unemployment risk, this effect does not break the positive and proportional relation between permanent earnings and the standard deviation of future earnings.

present those results in Appendix B.6.

What people are asked about in the SCE is the annual earnings they expect four months from now. The fact that the question is about annual earnings at a short horizon, four months from now, in March of the next year, is not per se a problem. Indeed, what I aim to measure is the level of permanent earnings that people have at a given point in time, not the change they experience over a year or a period of time longer than four months. Furthermore, people do not report the same values for current annual earnings and expected future annual earnings in March of next year (see Table 5 in Appendix B.3), so the horizon is long enough for them to expect some change. Current and future expected annual earnings also affect the MPC differently (see Table 3 in section 5).

Rozsypal and Schlafmann (2023) find that people are partly overestimating the persistence ρ of their earnings. I use self-reported expectations, so the ρ I should use is the one that people believe. Since I set ρ equal to one, this overestimation of the persistence is in line with my specification. Balleer, Duernecker, Forstner, and Goensch (2021) document an optimistic bias among SCE respondents with a low level of education. The subjective expectations of non-college graduates about their future labor market transitions are on average overoptimistic. In contrast, college graduates have rather precise beliefs. Since the bias strongly correlates with educational achievement, the education dummies that I use when detrending expected annual earnings should capture it.

4 Permanent earnings and the MPC in survey data

4.1 Specification and results

Specification. To measure the influence of permanent earnings on people's MPC, I estimate the following reduced-form specification:

$$MPC_t^i = a_1 + a_2 perm_t^i(1 + b_2 hh\ size_t^i) + a_3 wealth\ cat_t^i(1 + b_3 hh\ size_t^i) \quad (4.1)$$

$$+ a_4 hh\ inc_t^i(1 + b_4 hh\ size_t^i) + a_5 hh\ size_t^i + a_6 (wealth\ cat_t^i * hh\ inc_t^i) \quad (4.2)$$

$$+ a_7 spouse\ lf_t^i + a_8 age_t + \xi_t^i.$$

The term MPC_t^i is the reported MPC out of a hypothetical shock of respondent i at period t . The term $perm_t^i$ is the permanent component of the earnings of the respondent, built as described above. The term $hh\ size_t^i$ is a vector of dummies for the number of household members and for whether or not there are children in the household. The term $wealth\ cat_t^i$ is a vector of dummies for 14 categories of household non-housing wealth. The term $hh\ inc_t^i$ is a vector of dummies for 11

categories of combined pre-tax income of all household members. One reason why I control for household income, besides the fact it is likely to influence the MPC, is that the MPC question asks about a shock proportional to household income. By controlling for household income, I thus control for shock size. I do this because my theoretical result, and a large part of the literature, is about the determinants of the response to a shock of the same size for all. The term $spouse\ l f_t^i$ is a vector of dummies for the labor force status of the spouse, including the absence of any spouse, a full-time working spouse, a part-time working spouse, a retired spouse, or a home-maker spouse. The term age_t^i is a vector of dummies for six age categories. The term ξ_t^i is the noise.

I also run a second specification with year and demographic controls, which include the educational attainment and the willingness to take risks of the respondent, the state in which the household resides, and whether or not the household has a budget or plan for their monthly spending and saving.

Implementation. I estimate the specifications described by (4.1) with a linear regression.

	MPC loss	MPC gain	MPC loss	MPC gain
Permanent earnings (in \$10,000)	0.012*** (0.004)	0.007** (0.004)	0.013*** (0.004)	0.010** (0.004)
One s.d. change in permanent earnings	0.040***	0.025**	0.043***	0.033**
5th to 95th percentile of permanent earnings	0.125***	0.077**	0.134***	0.103**
1st to 11th category of wealth (\approx 5th to 95th percentile)	-0.191***	-0.239***	-0.200***	-0.279***
Average MPC	0.788	0.462	0.788	0.462
Full demographic controls	No	No	Yes	Yes
Observations	1,100	1,100	1,099	1,099
R^2	0.249	0.374	0.302	0.418

Robust standard errors in parentheses. * at 10%, ** at 5%, *** at 1%.

Table 1: Average marginal effect of permanent earnings on the MPC

The effect of permanent earnings on the MPC. Table 1 presents the results obtained from the estimation of specification (4.1). The first two columns present the results based on the regression without all demographic controls. In the first column, the dependent variable is the MPC out of an income loss. The first line shows that the point estimate of the average marginal effect of permanent earnings is 0.012, significant at the 1% level.¹³ This means that, everything else being

¹³Expressing it with the coefficients in equation 4.1, this average marginal effect corresponds $a_2(1 + b_2\overline{hh\ size}_t^i)$,

equal and at the average household size, a \$10,000 higher level of permanent earnings of a head associates with a 0.012 higher MPC out of a loss: in response to a loss people would cut their consumption by an extra 1.2% of the loss. This is in line with the theoretical prediction of the standard model that I uncover: at a higher level of permanent earnings, people respond more to one-time shocks. The next three lines expand on the magnitude of the effect. The point estimate of 0.012 implies that a one-standard deviation increase in permanent earnings raises the MPC out of a loss by 0.040. Such an increase represents 5% of the average value of the MPC in the sample, which is 0.788. It also implies that moving from the 5th to the 95th percentile of the permanent earnings distribution (while keeping the variables with which it interacts in equation (4.1) fixed and equal to their average value) raises the MPC out of a loss by 0.125. This magnitude of the effect is comparable to that of liquid wealth. I find that moving from the 1st wealth category, corresponding to the 5th percentile of wealth, to the 11th category, corresponding to the 95th percentile, (while keeping the variables with which it interacts in equation (4.1) fixed and equal to their average value) reduces the MPC out of a loss by 0.191. Thus, the position on the permanent earnings distribution can offset a substantial part of the effect of the position on the non-housing wealth distribution.

The second column shows that these results still hold true when considering the MPC out of an income gain. This is important because the theory predicts that permanent earnings raise the sensitivity of consumption to both types of shocks. The point estimate of the average effect is 0.007, significant at the 5% level. Looking at the magnitude of the effect through additional statistics, the result implies that a one standard deviation increase in permanent earnings raises the MPC out of a gain by 0.025. This is 5% of the average value of this MPC, which is 0.462. Moving from the 5th to the 95th percentile of the permanent earnings distribution raises the MPC out of a gain by 0.077, while moving from the wealth category equivalent reduces this MPC by 0.239. Thus, the extent to which moving along the permanent earnings distribution offsets the negative effect of moving along the wealth distribution is one third for the MPC out of a gain, instead of two thirds for the MPC out of a loss.

The next two columns show that the results are similar and slightly larger when I add demographic controls including education, which predicts overoptimism about future earnings, or willingness to take risks, which captures risk-aversion.

One policy implication is that, in order to maximize the consumption response, targeting stimulus checks to people based on their income is not the most effective. This is because income includes components that both raise and reduce the MPC. The ideal would be to be able to observe permanent income and wealth: people in the bottom two wealth categories and above the median

with $\overline{hh\ size}_t^i$ the average values of the household size dummies in the sample.

level of permanent earnings have an average MPC out of a gain of 0.66. This is a substantial increase compared to the average MPC of 0.47 in the whole sample and 0.50 in the bottom 10th percentile of the earnings distribution. Absent this possibility, targeting people below age 45 in the bottom 10th percentile of the earnings distribution can raise the average MPC out of a gain to 0.53.

Explaining fact (i): some people with high levels of liquid wealth still have a non-zero MPC, and the effect of liquid wealth on the MPC is modest. The first stylized fact that motivates the analysis is that some people with medium and high levels of liquid wealth still have a non-zero MPC, resulting in a modest effect of liquid wealth on the MPC. The MPC that this literature relies on is the MPC out of a gain so that is the one I focus on. However, what I document also holds for the MPC out of a loss in the SCE. I observe that, the average MPC out of a gain for households in the 6th category of non-housing wealth, who hold the medium amounts of \$10,000 to \$20,000, is still 0.44; the average MPC out of a gain of households in the top two, 13th and 14th, categories is 0.27—I bundle the two together since there is only 18 observations in the top one. Also, the difference between the average MPC out of gain in the bottom and top two wealth categories is only 0.35—from 0.62 to 0.27.

The mechanism that I uncover can explain this stylized fact. First, I find that the permanent component of earnings affects the MPC in the opposite direction as liquid wealth does. Second, when I compare the size of this effect with that of liquid wealth, I find that they have the same order of magnitude. Because wealth and permanent earnings are partly correlated, people with substantial liquid wealth face more earnings risk and their MPC can still be high.

Explaining fact (ii): conditional on wealth, people with higher current earnings do not have a significantly lower MPC. The second stylized fact that motivates the analysis is that, conditional on wealth, people with higher current earnings do not respond significantly less to one-time income shocks. This is surprising because one might have expected the effect of current earnings to be similar to that of liquid wealth, since both provide additional immediately available resources. I confirm this stylized fact in the SCE: when I estimate a version of specification (4.1) in which I substitute permanent earnings with total current earnings, the average marginal effect of total earnings becomes non-significant and small. The point estimates are small: equal to 0.007 (standard error of 0.005) on the MPC out of a loss, and 0.001 (standard error of 0.005) on the MPC out of a gain. I reproduce those results in section 5, Table 3, for comparison with the numerical results. These findings are in contrast with the effect of liquid wealth on the MPC, which is significant and negative.

The mechanism that I uncover can account for this. While an increase in current earnings does provide additional immediately available resources, it also increases future not-yet-realized resources and their variance when the increase is coming from the permanent component of current earnings. Overall, because current earnings is made of both a permanent component, which raises the MPC, and a transitory component, which reduces it, its average effect on the MPC can be small and not significant.

4.2 Robustness: similar results from other methods and specifications

Bootstrapping. My measure of permanent earnings is detrended expected future annual earnings conditional on future employment. Although the practice of working with detrended variables is commonplace in the consumption literature, the first step regression that I run for detrending can have some extra variability that is not reflected in the standard errors of the main estimation. To account for this extra variability, I recompute the standard errors with a bootstrap procedure that includes the first step in a bootstrapping loop of 500 iterations. This introduces little change in the standard errors: they increase from 0.004 to 0.005. The detailed results are in Table 9 Appendix C.1.

Broader notion of permanent earnings: an almost direct regression of expected future earnings on the MPC. In the baseline specification, I am excluding the variations in the fixed effect component of earnings that are explained by demographic variables. This means that I am excluding, for instance, differences in the permanent component of earnings driven by differences in education. To include these effects, I build permanent earnings by only detrending expected future earnings conditional future employment from the effect of the gender of the respondents—I still remove the effect of gender because studies have shown that the gender is strongly correlated with the MPC (see Boehm, Fize, and Jaravel (2023)). The effect of this broader measure of permanent earnings remains significant and the magnitudes are similar. Table 10 in Appendix C.2 presents the results.

Generalized specifications. The specification described by (4.1) imposes a linear relation between the MPC and permanent earnings.¹⁴ I examine what happens when I allow for higher order interactions. I let permanent earnings interact with household income and with wealth. Table 11 in Appendix C.3 presents the results. The average partial effect of permanent earnings on the MPCs

¹⁴Note that it does not impose a linear relationship between consumption and permanent earnings, but only between the MPC, that is, the partial effect of liquid wealth on consumption, and permanent earnings.

remains significant and positive.

Using consumption data rather than hypothetical MPCs to examine the effect of earnings on MPCs. As a robustness exercise, I consider a different specification that relies on questions about realized consumption rather than on questions about the response to hypothetical shocks. In this specification, I estimate the interaction between the effect of permanent earnings and the effect of non-housing wealth on consumption. Indeed, the effect of non-housing wealth on consumption measures a form of MPC. Thus, their interaction is a proxy for the influence of permanent earnings on this MPC. Note that what it captures still differs from the effect of permanent earnings on a true MPC for at least three reasons. First, the correlation between a change in liquid wealth and consumption is not exactly a MPC out of an unexpected shock: liquid wealth changes are not necessarily exogenous and might reflect a response to other events also affecting consumption directly. That is why people rely on natural experiments rather than on regressions of consumption over liquid wealth to measure MPCs. Second, the consumption level is indirectly recovered from other variables thus obtained for only a fraction of the sample. It also covers only typical consumption, excluding large infrequent purchases, while the hypothetical questions covers total consumption as it includes any spending. Third, the variations in liquid wealth are coming from variations in a categorical variables, thus less precise than if the variable had initially been continuous. With these limitations in mind, I find that the interaction between the effect of liquid wealth and of permanent earnings on typical consumption is significant and positive. This result is in line with what I obtain using hypothetical MPCs: everything else being equal, the consumption of people with more permanent earnings is more sensitive to changes in liquid wealth. Quantitatively, the magnitude of the effect is larger but not by a different order of magnitude: a \$10,000 increase in permanent earnings raises the effect of liquid wealth on consumption by 0.035, significant at the 1% level. I detail the specification and results in Appendix C.4.

Effect of the ratio of wealth to permanent earnings. The model predicts that what matters for the MPC is only the ratio of liquid wealth to permanent earnings. Permanent earnings should not impact the MPC beyond their effect on the ratio. I test for this. One difficulty is that I only observe the household's wealth category, not their exact wealth. To still examine this point, I convert the fourteen discrete wealth category dummies into a continuous wealth measure by attributing to each household the level of wealth corresponding to the lowest level of their category, and attributing zero wealth to those in the lowest non-housing wealth category. I divide this continuous measure of wealth by the level of permanent earnings. I discretize this ratio into eight categories. I substi-

tute my wealth dummies with the discretized ratio dummies in the estimation of (4.1). The effect of permanent earnings is no longer significant and the point estimate becomes much smaller. This is not driven by the de-discretization and re-discretization: when I convert the fourteen discrete wealth category dummies into a continuous wealth measure, re-discretize it into eight categories, and substitute my wealth dummies with these re-discretized wealth dummies, the effect of permanent earnings remains as significant and as large as in my baseline. I present those results in Appendix C.5.

5 Permanent earnings and the MPC in simulated data

To understand whether life-cycle models are not only qualitatively but also quantitatively consistent with my empirical findings, I run numerical simulations. Indeed, precautionary effects typically generate behaviors that are qualitatively consistent with a number of puzzling consumption stylized facts but, quantitatively, their effects often ends up being too small to really account for them.

5.1 Model and calibration

Consumers' maximization problem. I simulate and calibrate a standard incomplete market model that mimics the situation of US households. A household is made of one individual solving a similar consumption maximization problem as the one I describe in Section 2. A period is a year, since this is the timespan that people are asked to consider in the survey. The period utility $u(\cdot)$ is a log-utility function. There is an age-specific extra discount factor equals to 0.985, that multiplies the discount factor at every age from 49 years old on. This is to match the hump-shaped pattern of consumption over the life-cycle, which Attanasio, Banks, Meghir, and Weber (1999) and Attanasio (1999) document. It should capture that people are done paying for some expenses that are life-cycle specific (e.g. paying for children's expenses and education, or doing more home-production in retirement, as documented in Aguiar and Hurst (2005), Aguiar and Hurst (2007), and Hurd and Rohwedder (2013)). I choose age 49 because that is the shifting point in the hump-shaped patterns documented in Attanasio, Banks, Meghir, and Weber (1999) and Attanasio (1999).¹⁵

Wealth. I assume that wealth in the model represents net risk-free liquid wealth. This means assuming that people may have illiquid wealth on the side, but that they do not use it to smooth

¹⁵See Figure 1 in Attanasio, Banks, Meghir, and Weber (1999) and Figure 4 in Attanasio (1999).

consumption, following the insight in Kaplan and Violante (2014).¹⁶ Kaplan and Violante (2022) further show that a one-asset model that matches the level of liquid wealth that people hold, rather than their total wealth, can generate MPCs consistent with empirical evidence and as large as in a two-assets model that explicitly has both liquid and illiquid wealth. Thus, while a model with two assets makes it possible to match evidence on both consumption and wealth, the one-asset model seems to be an adequate simplification when the objective is to model consumption.

The yearly interest rate on the liquid asset is constant and set to $r = 0.01$, to match the low real interest rate on liquid holdings over the period 2015-2018.¹⁷

Discount factor. I calibrate internally the baseline discount factor β that applies before age 49 and is multiplied by 0.985 after age 49. I set it so that the mean value of liquid wealth in the population equals 20% of the mean annual earnings in the population. This share of 20% is the same calibration target as in Kaplan and Violante (2022), chosen to be consistent with the 2019 Survey of Consumer Finances (SCF). The mean earnings in my baseline model is \$61,638, so this implies a mean liquid wealth of \$12,327 (matched +/-1%). The discount factor before age 49 that yields it is $\beta = 0.969$. This implies that the discount factor after age 49 is $0.969 * 0.985 = 0.954$

Borrowing limit. In addition to the period budget constraints, people face a borrowing limit on how much they can borrow for consumption purposes. In the baseline calibration, I fix it at a maximum consumption debt of \$5,500 (in \$2014). This is coming from the SCF data about the credit card balance still owned (question x413). The top 90th percentile of balance still owned is \$4,939 in the 2016 survey and \$6,121 in the 2019 survey (both deflated and expressed in \$2014). Since my period is between the end of 2015 and the end of 2018, I take their average as my target and set the borrowing constraint at \$5,500.

Lifespan and survival probabilities. People enter the labor market at age 25. They retire at age 62. After retirement, people have a non-zero probability to die at each period from age 62 to age

¹⁶The model is for instance equivalent to one in which people would either rent their whole life or start their working life with a house and an exogenous amount of mortgage they have to repay at each period. Their mortgage payment correspond to their expenses in housing services, which I take into account in the data. They never sell the house. When they die, they pass on a fraction of the house to each of their children, who use a money plus a mortgage to buy their own house. The children are then in the same situation as their parents at the beginning of their working life, with a house and a mortgage to repay.

¹⁷The average 10-Year Real Interest Rate in the US over 2015-2018 is 0.0072 (see *Federal Reserve Bank of Cleveland, 10-Year Real Interest Rate [REINTRATREARAT10Y]* (n.d.)). Incidentally, because the discount factor β is set to match the empirical level of liquid wealth, changing the interest rate leads to an adjustment in the internally calibrated β and has little impact on the simulation results.

91. I obtain the survival probabilities from the life tables of the National Center for Health Statistics.¹⁸ If still alive at age 91, a household dies with certainty at age 92.

Earnings. The earnings that people receive at each period follow exactly the parametric process (5) proposed in Guvenen, Karahan, Ozkan, and Song (2021). It is the same as the process I consider in the empirical section, except that in the empirical section I did not have to take a stand on the distribution of the shocks and on the functional form of the probability of non-employment. Here, I follow the distributions and functional forms of Guvenen, Karahan, Ozkan, and Song (2021):

$$\text{Annual earnings: } y_t^i = \underbrace{(1 - v_t^i)}_{\text{Employment status}} \underbrace{e^{\alpha^i}}_{\text{Fixed effect}} \underbrace{e^{p_t^i}}_{\text{Highly persistent}} \underbrace{e^{\varepsilon_t^i}}_{\text{Transitory}} \underbrace{e^{g(t)}}_{\text{Age trend}} \quad (5.1)$$

$$\text{Persistent component: } e^{p_t^i} = (e^{p_{t-1}^i})^\rho e^{\eta_t^i}, \quad (5.2)$$

$$\text{Nonemployment: } v_t^i \sim \begin{cases} 0 \text{ (employment) with prob. } 1 - p_{v_{t-1}^i}, \\ 1 \text{ (nonemployment) with prob. } p_{v_{t-1}^i}, \end{cases} \quad (5.3)$$

$$\text{Prob. of nonempl.: } p_v(t, e^{p_t^i}) = \frac{e^{\xi_t^i}}{1 + e^{\xi_t^i}} \text{ where } \xi_t^i \equiv a_v + b_v t + c_v p_t^i + d_v t p_t^i, \quad (5.4)$$

$$\text{Persistent innovation: } \eta_t^i \sim \begin{cases} \mathcal{N}(\mu_{\eta,1}, \sigma_{\eta,1}^2) \text{ with prob. } p_\eta, \\ \mathcal{N}(\mu_{\eta,2}, \sigma_{\eta,2}^2) \text{ with prob. } (1 - p_\eta), \end{cases} \quad (5.5)$$

$$\text{Transitory innovation: } \varepsilon_t^i \sim \begin{cases} \mathcal{N}(\mu_{\varepsilon,1}, \sigma_{\varepsilon,1}^2) \text{ with prob. } p_\varepsilon, \\ \mathcal{N}(\mu_{\varepsilon,2}, \sigma_{\varepsilon,2}^2) \text{ with prob. } (1 - p_\varepsilon), \end{cases} \quad (5.6)$$

$$\text{Fixed effect: } \alpha^i \sim \mathcal{N}(0, \sigma_\alpha^2) \quad (5.7)$$

$$\text{Initial persistent: } p_0^i \sim \mathcal{N}(0, \sigma_{p0}^2). \quad (5.8)$$

I set the parameters of this process equal to the estimates of Guvenen, Karahan, Ozkan, and Song (2021), summarized in Appendix D.1 of this paper and taken from Table IV of their paper and Table D.III of their online appendix.

Taxes and social security income. People pay taxes according to the nonlinear tax function of Gouveia and Strauss (1994), $\text{tax}(y_t^i) = \tau^b (y_t^i - ((y_t^i)^{-\tau^p} + \tau^s)^{-1/\tau^p})$ parametrized with $\tau^b = 258$, $\tau^p = 0.768$, $\tau^s = 2.0e - 4$ as in Kaplan and Violante (2010).¹⁹

¹⁸See https://www.cdc.gov/nchs/products/life_tables.htm (accessed May 2024).

¹⁹Contrary to Kaplan and Violante (2010) who model net income and use the inverse of the tax function to recover gross income, here, what I model with the Guvenen, Karahan, Ozkan, and Song (2021) process is pre-tax earnings and I use the tax function to recover net earnings.

After retirement, people stop paying taxes and receive a social security income that is a deterministic function of their past income. More precisely, up to a given bend point, this social security income is equal to 90 percent of average past earnings. It is 32 percent from this first bend point to a second bend point, and 15 percent beyond that. The two bend points are set at 0.18 and 1.10 times the cross-sectional average gross earnings. This follows Kaplan and Violante (2010), who mimic the US legislation.

MPCs. To compute people’s MPCs, I simulate two alternative situations for every individual, on top of the one without shocks. In the first one, they are hit by a one-time negative shock and their beginning of period wealth decreases by 10% of their current yearly income. In the second one, they are hit by a one-time positive shock and their beginning of period wealth increases by 10% of their current yearly income. These are the shocks that people are asked about in the survey question. The shocks occur at random, once in the life-time, between age 26 and age 55. For a given individual, the two shocks occur at the same point in the lifetime. I compute the MPCs as the consumption difference with and without the shock over the size of the shock.

5.2 Simulations

Method. I simulate an artificial panel of 5,000 consumers, and I solve the model using the method of endogenous grid points developed in Carroll (2006).²⁰

Price harmonization. In the simulations, the income process is calibrated with the parameters estimated by Guvenen, Karahan, Ozkan, and Song (2021). Their estimation is based on data deflated and expressed in 2010\$ value. I thus simulate the model in 2010\$ and convert the simulated values to 2014\$. The borrowing limit target that I set accounts for this price harmonization: to impose a borrowing limit of \$5,500 in 2014\$, I set it at \$5,058 in the numerical simulations.

Building permanent earnings. I directly observe the fixed effect α^i and the highly persistent component of log-earnings p_t^i . I normalize their product in the same way I do with survey data: I regress it over the year dummies (or equivalently the age dummies since the two coincide in the simulations), take the exponential of the residual, and multiply it with the exponential of average

²⁰The number of grid points is as follows: the grid for wealth has 150 exponentially spaced grid points; the grid for the highly persistent component of earnings is age-varying and at each age has 35 equally spaced points; the grid for the transitory shock has 11 equally spaced points; the grid for the fixed effect component of earnings has 9 equally spaced points; the grid for lifetime average earnings (used to compute retirement income) has 9 equally spaced points. Expanding the grid further does not change the results.

log-earnings among employed people.

Selection. As in the empirical analysis, I select individuals aged 25-55 and employed at the moment when they experience the transitory shock. I trim the top and bottom 1% of the permanent component of earnings. The calibration of the parameters is done before this trimming.

	Survey data (SCE)	Simulated data
Average earnings	63,592	61,638
Average permanent earnings	60,207	57,570
Correlation permanent earnings/wealth categories	0.317	0.238
Share of people at the constraint	.	0.333
Observations	1,099	3,308

Table 2: Model fit

Wealth and earnings comparison in the simulated and survey data. How close is the model to the data on the earnings and wealth dimensions? The average earnings generated by the process in Guvenen, Karahan, Ozkan, and Song (2021) is close to, though a little below, the average earnings of the respondents in the SCE. The difference might be due to the fact that respondents in the SCE are household heads (who contribute to the rent or own the house), who might earn a higher wage than non-heads. As a consequence, the average permanent earnings is also a little lower in the simulations. The correlation between permanent earnings and the non-housing wealth categories is positive but not too large in both the survey and the simulated data. It is a little higher in the survey data. Finally, in the simulated data, a little less than one third of the people have their liquid wealth at the minimum possible level of $-\$5,500$. This share is similar to the estimated share of hand-to-mouth people in the population, that is, people with very low levels of liquid wealth (with or without illiquid wealth on the side): the baseline share of hand-to-mouth in the seminal paper of Kaplan, Violante, and Weidner (2014) is 0.312, with a range going from 0.220 to 0.503.

5.3 The effect of permanent earnings on the MPCs in simulated data

Specification. With these simulations data, I estimate a specification close to (4.1), the one that I estimate in survey data. For the wealth dummy variables, I discretize the liquid wealth variable with the same thresholds as in the survey question. The only demographic dummy variables are the age dummy variables. I create one age category dummy for each age level. Because all households

are single, the household income coincides with total earnings. I discretize it in eleven categories of the same size to obtain the household income categories. The equation that I estimate is then:

$$MPC_t^i = a_1 + a_2 perm_t^i + a_3 wealth_t^i + a_4 hh inc_t^i + a_5 (wealth cat_t^i * hh inc_t^i) + a_6 age_t + \xi_t^i.$$

	Survey (SCE)		Simulations (baseline)	
	MPC Loss	MPC Gain	MPC Loss	MPC Gain
Permanent earnings (in \$10,000)	0.013*** (0.004)	0.010** (0.004)	0.011 .	0.009 .
Average MPC	0.788	0.462	0.616	0.566
R^2	0.302	0.418	0.867	0.803
Observations	1,099	1,099	3,266	3,266
Total earnings (in \$10,000)	0.007 (0.005)	0.001 (0.005)	0.001 .	0.001 .
R^2	0.298	0.411	0.868	0.804
Observations	1,083	1,083	3,308	3,308

Robust standard errors in parentheses. * at 10%, ** at 5%, *** at 1%.

Table 3: Effect of persistent earnings on the MPC in survey data and in simulations

Results. Table 3 presents a comparison of the effect of permanent earnings on the MPC in the survey data and in the simulated data. The first two columns of the top panel are a reminder of the results I obtain in the survey data. The third and fourth columns of the top panel present the effect of permanent earnings on the MPC obtained from the estimation of equation (5.9). In the simulated data, a \$10,000 increase in permanent earnings raises the MPC out of a loss by 0.011, close to the increase by 0.013 that I estimate in survey data. Similarly, a \$10,000 increase in permanent earnings raises the MPC out of a gain by 0.009, also close to the survey data estimate of 0.010.

Incidentally, the R^2 s of the estimation regression are large in the simulated data, which confirms that the simple linear relation between the MPC and persistent earnings that I estimate captures a large share of the MPC fluctuations that are generated by a life-cycle model.

High MPCs Importantly, this model is able to generate large MPCs. The average MPC out of a gain is 0.616 and the average average MPC out of a loss is 0.566. Compared to the MPCs of 0.788 and 0.462 reported in survey data, the levels are in the same range. The only difference between this model and the ones typically used for numerical simulations is the more realistic

earnings process that I use. This means that simply including an earnings process that incorporates more earnings risk can substantially increase the average MPC that the model generates. This matters because generating such high MPCs out of a life-cycle model has proved difficult in the past. When Kaplan and Violante (2022) review different models of this style, the average annual MPC out of a positive shock implied is rarely above 0.45. Some papers rely on extensions that are internally calibrated to raise the average MPC of the model. Here, the change that I introduce is in an externally observable dimension.

Regarding the asymmetry between the two MPCs, existing numerical simulations focus on the MPC out of a gain. I observe both. The MPC out of a loss is higher than the MPC out of gain both in the survey data and in the simulated data, but the gap is more pronounced in the survey data. That is one thing that the model partly misses.




Figure 2: Average MPC out of a gain by liquid wealth in survey data (left) and simulations of the baseline model (right) in the bottom 90% of the liquid wealth distribution.

Stylized fact (i) in the simulated data The numerical simulations can reproduce the stylized fact (i), that people with non-zero liquid wealth still respond substantially to a one-time shock. To show this, I plot the evolution of the MPC out of a gain with wealth. I select out the top 10% of the survey and simulated data. In the survey data, this corresponds to dropping the top 4 wealth categories. I do this because the model does not include capital income and is not perfectly adequate for the behavior of the top 10%. The left panel in Figure 2 shows the average MPC out of a gain in the 10 bottom non-housing wealth category of the survey data. The right panel presents the average MPC in 10 same-sized wealth categories of the bottom 90% of the wealth distribution.

In both survey and simulated data, the average MPC decreases with non-housing wealth or liquid wealth, but still remains substantially above zero even in the middle-top of the wealth distribution. The slope is steeper in the simulations than in the survey data for the low wealth categories.

However, in both cases, the MPC decreases less steeply after the 5th category and plateaus around 0.2-0.3 at the top.

Stylized fact (ii) in the simulated data The numerical simulations can also reproduce the absence of a negative correlation between total current earnings and MPC. The last line in Table 3 presents the effect of total earnings, when estimating equation (4.1) and equation (5.9) with total earnings instead of permanent earnings. The first two columns show that in the survey data the effect of total earnings on the MPCs is not statistically significant, with a small positive point estimate. The third and fourth column show that the same result holds in the numerical simulations: the effect of total earnings on the MPC is close to zero, with a positive point estimate. This suggests that, as in the survey data, the positive effect of permanent earnings on the MPCs offsets the negative effect of the transitory component, so the effect of a change in total earnings is small.

Model Variations I show that the results barely change when I remove the extra discount factor after age 49 or extend the borrowing limit to \$30,000. However, the average MPCs drop below 0.3 when I remove entirely the borrowing limit. Both the average MPC and the effect of permanent earnings decrease largely when I calibrate the discount factor β to match the total wealth that people typically hold. Finally, I consider a model similar to the one I present in the theoretical section and with a typical calibration of the earnings process, that is, without any borrowing constraint, without social security income linked to past earnings, and with a simple transitory-persistent process with normal shocks and without non-employment shocks. In this model, the effect of permanent earnings on the MPC is still positive but much smaller. The average MPCs are much smaller as well. I present those results in Appendix D.2.

5.4 The role of the rich earnings process

My simulations are able to generate high MPCs and to account for the two stylized facts that motivate the analysis. I show that the richer earnings process is key to match those empirical observations, and I examine the features of the process that are the most important to do so.

Alternative earnings process I simulate alternative models in which the earnings process is different. In the first one, I shift to a simple transitory-persistent process with normally distributed shocks. This is the typical process used in numerical simulations of a standard incomplete market

model

$$\text{Annual earnings: } y_t^i = \underbrace{e^{\alpha^i}}_{\text{Fixed effect}} \underbrace{e^{p_t^i}}_{\text{Highly persistent}} \underbrace{e^{\varepsilon_t^i}}_{\text{Transitory}} \quad (5.9)$$

$$\text{Persistent component: } e^{p_t^i} = (e^{p_{t-1}^i})^\rho e^{\eta_t^i}, \quad (5.10)$$

$$\text{Persistent innovation: } \eta_t^i \sim \mathcal{N}(0, \sigma_\eta^2) \quad (5.11)$$

$$\text{Transitory innovation: } \varepsilon_t^i \sim \mathcal{N}(0, \sigma_\varepsilon^2) \quad (5.12)$$

$$\text{Fixed effect: } \alpha^i \sim \mathcal{N}(\mu_\alpha, \sigma_\alpha^2). \quad (5.13)$$

$$\text{Initial persistent: } p_0^i \sim \mathcal{N}(0, \sigma_{p0}^2). \quad (5.14)$$

I keep the calibration of the alternative process very close to the one of the rich baseline process. The persistence of the highly persistent component is the same $\rho = 0.991$. In the baseline process, people draw the persistent and transitory innovations from mixtures of normal distributions: with a high probability they draw from a normal distribution with a typical variance, but with a small probability they can draw from a normal distribution with a high variance. Here, the shocks are only drawn from one normal distribution each. Their variances are the same as the variances of the most probable distribution in the baseline. The mean of the distributions are zero. The variance of the fixed effect is the same as in the baseline. The mean of the fixed effect is set to match the same average annual earnings as in the baseline. I recalibrate the discount factor β so the average liquid wealth remains equal to 20% of average annual earnings.

In each of the other alternative models, I add back one element of the baseline earnings process, to see the impact of this element alone. These elements are: having people draw persistent innovations from a mixture of normal distributions, having people draw transitory innovations from mixture of normal distributions, having non-employment shocks, and having an upward sloping quadratic time trend. In each case, I adjust the mean of the fixed effect component to match the same average annual earnings as in the baseline. I recalibrate the discount factor so the average liquid wealth remains equal to 20% of average annual earnings.

The selection is the same as in the baseline. When there are no non-employment shocks and no persistent shocks drawn from high-variance distributions, the sample is larger since there are no non-employed individuals and less very low-earnings individuals to drop.

The effect of permanent earnings with and without the rich earnings process The first line of the top and bottom panels of Table 4 presents the average effect of permanent earnings on the MPCs in each alternative model. The first two columns reproduce the results of the baseline simulations.

	Baseline		Simple earnings		Mixture - persistent	
	Loss	Gain	Loss	Gain	Loss	Gain
Permanent earnings (in \$10,000)	0.011	0.009	-0.029	-0.018	-0.004	-0.005
Average MPC	0.616	0.566	0.597	0.173	0.557	0.290
Total earnings (in \$10,000)	0.001	0.001	-0.026	-0.015	-0.001	0.000
Observations	3,308	3,308	4,939	4,939	4,893	4,893
	Mixture - transitory		Non-employment		Quadratic trend	
	Loss	Gain	Loss	Gain	Loss	Gain
Permanent earnings (in \$10,000)	0.003	0.000	0.005	-0.003	0.027	0.024
Average MPC	0.579	0.169	0.531	0.488	0.611	0.400
Total earnings (in \$10,000)	-0.022	-0.012	0.002	-0.003	-0.013	-0.011
Observations	4,935	4,935	4,154	4,154	4,932	4,932

Robust standard errors in parentheses. * at 10%, ** at 5%, *** at 1%.

Table 4: Average MPCs in numerical simulations

The third and fourth columns show that, with the simple earnings process, permanent earnings reduce the MPC. This is because of the exogenous borrowing constraint. Recall that without an upward sloping life-cycle trend to earnings, an increase in permanent earnings relaxes the constraint and reduces the MPC. With the same simple earnings process but no exogenous borrowing constraint, the simulations of the 'Theoretical section' model, in Appendix D.2 shows that permanent earnings has a positive effect, not a negative one, on the MPC. The other columns show that most of the elements of the rich earnings process reduce this negative effect. Having shocks drawn from normal distributions or drawing extra non-employment shocks makes the effect of permanent earnings less negative, or even positive. This is presumably because they increase the earnings risk that people face and thus the impact that a higher permanent earnings that multiplies this risk has on their precautionary motive. The quadratic trend also has a strong impact on the effect of permanent earnings on the MPC. This is because introducing an upward sloping life-cycle trend to earnings shifts the direction of the impact of permanent earnings on the liquidity constraint: without any trend, an increase in permanent earnings relaxes the constraint, while with an upward sloping trend an increase in permanent earnings can strengthen the constraint. Note that saving (which can mean not borrowing as much as one would) for a precautionary motive and saving (which can mean not borrowing as much as one would) because of a liquidity constraint are substitute. As a result, in the baseline specification with more risk and a stronger precautionary motive, the effect of permanent earnings on the MPC can be smaller than with a quadratic life-cycle earnings trend only.

High MPCs with and without the rich earnings process The second lines in the top and bottom panels of Table 4 presents the average MPCs in each set of simulations. When I switch from the baseline to the simple earnings process typically used, the MPC out of a negative shock remains relatively large, 0.597, close to the baseline and a little below the survey response. The MPC out of a positive shock, however, drops to 0.173, substantially smaller than in the baseline and in the survey data. This asymmetry between the MPC out of a positive and negative shocks is presumably due to the fact that, in this model, the reason why people respond to a one-time shock is mostly because of liquidity constraints. A large positive shock corresponding to 10% of yearly income can move people outside of the constrained region where they are no longer very responsive. This suggests that the rich earnings process is key in generating a large MPC out of a positive shock specifically. The other columns show that introducing non-employment shocks alone raises the MPC out of a gain from 0.173 to 0.488, much closer to the baseline value of 0.566. Introducing a quadratic trend also has a large effect on the MPCs. The third most effective element is having people draw persistent shocks from a mixture of normal distributions rather than from a single normal distribution. Having people draw transitory shocks from a mixture has little effect on the average MPCs in the model.

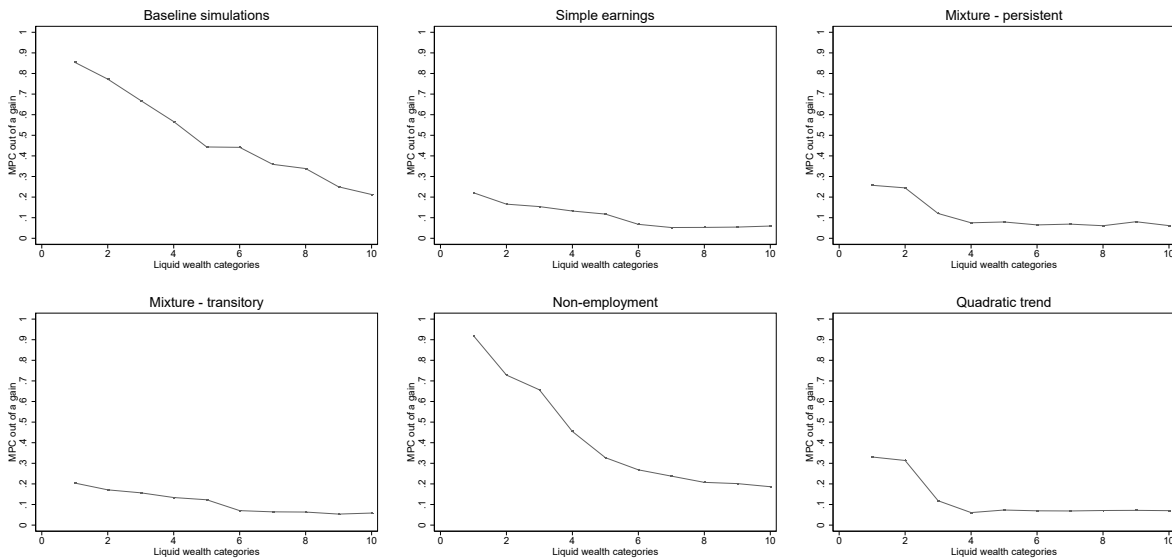


Figure 3: Average MPC out of a gain by measures of liquid wealth in survey data (left) and simulations of the baseline model (right).

Stylized fact (i) with and without the rich earnings process Figure 3 presents the evolution of the average MPC across 10 same-sized wealth category for the baseline as well as for the simula-

tions with the alternative earnings process. First, when shifting from the rich, baseline, earnings process (top left) to the simple earnings process (top middle), the average MPC out of a gain is much smaller and falls close to zero from the 6th category on. The other graphs presents the same evolution, except when the earnings process include non-employment shocks (bottom middle). With non-employment shocks only, the MPC plateaus from the 6th wealth category on and remains above 0.1-0.2. Even in that case, however, the MPC remains a little smaller at the top than in the baseline case. This suggests that the combination of non-employment shocks with the other elements also helps raise the MPC in the top of the wealth distribution.

Stylized fact (ii) with and without the rich earnings process The third lines in the top and bottom panels of Table 4 presents the effect of total earnings on the MPC. I obtain it by substituting permanent earnings with total earnings in the estimation of equation (5.9). As discussed, the baseline model is able to reproduce the empirical fact that the effect of total earnings is not negative but small and non-significant. This is no longer the case when I shift to the simple transitory-persistent earnings process typically used in simulations. The effect of total earnings becomes negative and larger in absolute value. This remains true, though to a lesser extent, when I include separately each of the component of the rich earnings process in the model. It suggests that the combination of the different elements is important to have a positive effect of permanent earnings on the MPC, and have that this effect dominates the negative effect of transitory earnings on the MPC.

6 Conclusion

In this paper, I establish the theoretical result that, in the standard life-cycle model used throughout macroeconomic studies, people with a higher permanent component of earnings have a higher MPC, everything else being equal. The result comes from precautionary behavior: people with a higher permanent component of earnings face a higher variance of their future earnings. Their consumption is relatively more constrained by uncertainty and they save more. A windfall gain relaxes this need for saving, and they consume more out of it than people who are not saving as much ex-ante. When the earnings process features a sufficiently upward-sloping life-cycle trend, people with a higher permanent earnings are also more likely to be constrained by an exogenous borrowing limit. This can further raise their MPC.

I examine the empirical validity of this theoretical prediction. I find that it holds true in the New York Fed Survey of Consumer Expectations. In this dataset, a one standard deviation increase in

permanent earnings associates with a 0.04 increase in the reported MPC out of a hypothetical one-time income shock. This is true both for the reported MPC out of a loss and for the reported MPC out of a gain. The effect of moving along the permanent earnings distribution is one-third to two-third as large as the effect of moving along the non-housing wealth distribution.

I then show that this empirical evidence is also quantitatively consistent with a standard consumption model calibrated to mimic the US economy: in numerical simulations of such a model, the effect of permanent earnings on the MPCs out of negative and positive shocks is as large as the one I measure in survey data. The MPC levels are also close to the ones I observe in survey data. The simulations can reproduce the empirical observations that the average MPC of people with substantial levels of wealth is still high (stylized fact (i)), and that current earnings do not have a negative effect on the MPCs (stylized fact (ii)). Incorporating the realistic and rich earnings process of Guvenen, Karahan, Ozkan, and Song (2021) is key to match quantitatively the empirical observations. It generates more earnings risk, which bolsters the precautionary motive thus the mechanism that I identify. Its upward-sloping life-cycle earnings trend also makes it more likely for people with a high permanent income to be constrained.

References

- Agarwal, S., Liu, C., and Souleles, N. (2007).** “The Reaction of Consumer Spending and Debt to Tax Rebates-Evidence from Consumer Credit Data.” *Journal of Political Economy* 115.6, pp. 986–1019.
- Aguiar, M. and Hurst, E. (2005).** “Consumption versus Expenditure.” *Journal of Political Economy* 113.5, pp. 919–948.
- Aguiar, M. and Hurst, E. (2007).** “Life-Cycle Prices and Production.” *American Economic Review* 97.5, pp. 1533–1559.
- Arellano, Manuel, Bonhomme, Stéphane, Vera, Micolle De, Hospido, Laura, and Wei, Siqi (Sept. 2021).** *Income Risk Inequality: Evidence from Spanish Administrative Records*. Working Papers 2136. Banco de España.
- Armantier, Olivier, Topa, Giorgio, Klaauw, Wilbert Van der, and Zafar, Basit (2017).** “An overview of the Survey of Consumer Expectations.” *Economic Policy Review* 23-2, pp. 51–72.
- Attanasio, O. (1999).** “Consumption.” *Handbook of Macroeconomics*. Ed. by Taylor, J. B. and Woodford, M. Vol. 1. Handbook of Macroeconomics. Elsevier. Chap. 11, pp. 741–812.
- Attanasio, O., Banks, J., Meghir, C., and Weber, G. (1999).** “Humps and Bumps in Lifetime Consumption.” *Journal of Business and Economic Statistics* 17.1, pp. 22–35.
- Attanasio, O., Kovacs, A., and Molnar, K. (2020).** “Euler Equations, Subjective Expectations and Income Shocks.” *Economica* 87.346, pp. 406–441.
- Aydin, D. (2019).** “Consumption Response to Credit Expansions: Evidence from Experimental Assignment of 45,307 Credit Lines.” *mimeo*.
- Bach, Laurent, Calvet, Laurent E., and Sodini, Paolo (2020).** “Rich Pickings? Risk, Return, and Skill in Household Wealth.” *American Economic Review* 110.9, pp. 2703–47.

- Baker, S. (2018).** “Debt and the Response to Household Income Shocks: Validation and Application of Linked Financial Account Data.” *Journal of Political Economy*. ISSN: 0022-3808.
- Balleer, Almut, Duernecker, Georg, Forstner, Susanne K., and Goensch, Johannes (2021).** *The Effects of Biased Labor Market Expectations on Consumption, Wealth Inequality, and Welfare*. CESifo Working Paper Series 9326. CESifo.
- Blundell, R., Pistaferri, L., and Preston, I. (2008).** “Consumption Inequality and Partial Insurance.” *American Economic Review* 98.5, pp. 1887–1921.
- Boehm, Johannes, Fize, Etienne, and Jaravel, Xavier (2023).** “Five Facts about MPCs: Evidence from a Randomized Experiment.” *Working Paper*.
- Borusyak, Kirill, Jaravel, Xavier, and Spiess, Jann (Apr. 2022).** *Revisiting Event Study Designs: Robust and Efficient Estimation*. CEPR Discussion Papers 17247. C.E.P.R. Discussion Papers.
- Boutros, Michael (Sept. 2022).** “Windfall Income Shocks with Finite Planning Horizons.” *Bank of Canada* 2022-40.
- Braxton, J. Carter, Herkenhoff, Kyle F, Rothbaum, Jonathan L, and Schmidt, Lawrence (2021).** *Changing Income Risk across the US Skill Distribution: Evidence from a Generalized Kalman Filter*. Working Paper 29567. National Bureau of Economic Research.
- Broda, Christian and Parker, Jonathan A. (2014).** “The Economic Stimulus Payments of 2008 and the aggregate demand for consumption.” *Journal of Monetary Economics* 68.S, pp. 20–36.
- Carroll, C. (2009).** “Precautionary saving and the marginal propensity to consume out of permanent income.” *Journal of Monetary Economics* 56.6, pp. 780–790.
- Carroll, C. and Kimball, M. (1996).** “On the Concavity of the Consumption Function.” *Econometrica* 64.4, pp. 981–92.
- Carroll, Christopher (2006).** “The method of endogenous gridpoints for solving dynamic stochastic optimization problems.” *Economics Letters* 91.3, pp. 312–320.
- Carroll, Christopher D., Holm, Martin B., and Kimball, Miles S. (2021).** “Liquidity constraints and precautionary saving.” *Journal of Economic Theory* 195, p. 105276. ISSN: 0022-0531.
- Commault, Jeanne (2024).** “On the Concavity of Consumption in the Permanent Component of Earnings.” working paper. see <https://www.jcommault.com/research/>.
- Crossley, Thomas F., Fisher, Paul, Levell, Peter, and Low, Hamish (2021).** “MPCs in an economic crisis: Spending, saving and private transfers.” *Journal of Public Economics Plus* 2, p. 100005. ISSN: 2666-5514.
- Fagereng, A., Holm, M., and Natvik, G. (2021).** “MPC Heterogeneity and Household Balance Sheets.” *American Economic Journal: Macroeconomics* 13.4, pp. 1–54.
- Fagereng, Andreas, Guiso, Luigi, Malacrino, Davide, and Pistaferri, Luigi (2020).** “Heterogeneity and Persistence in Returns to Wealth.” *Econometrica* 88.1, pp. 115–170.
- Federal Reserve Bank of Cleveland, 10-Year Real Interest Rate [REAINTRATREARAT10Y]* (n.d.). <https://fred.stlouisfed.org/series/REAINTRATREARAT10Y>, retrieved from FRED, Federal Reserve Bank of St. Louis. Accessed: May 31, 2024.
- Fuster, A., Kaplan, G., and Zafar, B. (2020).** “What Would You Do with \$500? Spending Responses to Gains, Losses, News, and Loans.” *The Review of Economic Studies*. rdaa076. ISSN: 0034-6527.
- Ganong, P., Jones, D., Noel, P., Farrell, D., Greig, F., and Wheat, C. (2020).** “Wealth, Race, and Consumption Smoothing of Typical Income Shocks.” Working Paper Series 27552.
- Gong, Liutang, Zhong, Ruquan, and Zou, Heng-fu (2012).** “On the concavity of the consumption function with the time varying discount rate.” *Economics Letters* 117.1, pp. 99–101.

- Gouveia, Miguel and Strauss, Robert P. (1994).** “Effective Federal Individual Tax Functions: An Exploratory Empirical Analysis.” *National Tax Journal* 47.2, pp. 317–39.
- Guvenen, Fatih, Karahan, Fatih, Ozkan, Serdar, and Song, Jae (2021).** “What Do Data on Millions of U.S. Workers Reveal About Lifecycle Earnings Dynamics?” *Econometrica* 89.5, pp. 2303–2339.
- Holm, Martin Blomhoff (2018).** “Consumption with liquidity constraints: An analytical characterization.” *Economics Letters* 167, pp. 40–42. ISSN: 0165-1765.
- Hurd, M. and Rohwedder, S. (2013).** “Heterogeneity in spending change at retirement.” *The Journal of the Economics of Ageing* 1, pp. 60–71.
- Jappelli, Tullio and Pistaferri, Luigi (2014).** “Fiscal Policy and MPC Heterogeneity.” *American Economic Journal: Macroeconomics* 6.4, pp. 107–36.
- Kaplan, G. and Violante, G. (2010).** “How Much Consumption Insurance beyond Self-Insurance?” *American Economic Journal: Macroeconomics* 2.4, pp. 53–87.
- Kaplan, G. and Violante, G. (2014).** “A Model of the Consumption Response to Fiscal Stimulus Payments.” *Econometrica* 82.4, pp. 1199–1239.
- Kaplan, G., Violante, G., and Weidner, J. (2014).** “The Wealthy Hand-to-Mouth.” *Brookings Papers on Economic Activity*. ISSN: 1533-4465.
- Kaplan, Greg and Violante, Giovanni L (2022).** *The Marginal Propensity to Consume in Heterogeneous Agent Models*. Working Paper 30013. National Bureau of Economic Research.
- Koşar, Gizem, Melcangi, Davide, Pilossoph, Laura, and Wiczer, David (June 2023).** *Stimulus through Insurance: The Marginal Propensity to Repay Debt*. Staff Reports 1065. Federal Reserve Bank of New York.
- Kueng, L. (2018).** “Excess Sensitivity of High-Income Consumers.” *The Quarterly Journal of Economics* 133.4, pp. 1693–1751. ISSN: 0033-5533.
- Lewis, Daniel J., Melcangi, Davide, and Pilossoph, Laura (2022).** *Latent Heterogeneity in the Marginal Propensity to Consume*. Staff Reports 1033. Federal Reserve Bank of New York.
- Meghir, C. and Pistaferri, L. (2004).** “Income Variance Dynamics and Heterogeneity.” *Econometrica* 72.1, pp. 1–32.
- Misra, K. and Surico, P. (2014).** “Consumption, Income Changes, and Heterogeneity: Evidence from Two Fiscal Stimulus Programs.” *American Economic Journal: Macroeconomics* 6.4, pp. 84–106.
- Orchard, Jacob, Ramey, Valerie A., and Wieland, Johannes F. (Aug. 2023).** *Micro MPCs and Macro Counterfactuals: The Case of the 2008 Rebates*. NBER Working Papers 31584. National Bureau of Economic Research, Inc.
- Parker, J., Souleles, N., Johnson, D., and McClelland, R. (2013).** “Consumer Spending and the Economic Stimulus Payments of 2008.” *American Economic Review* 103.6, pp. 2530–53.
- Parker, Jonathan A, Schild, Jake, Erhard, Laura, and Johnson, David (2022).** *Household Spending Responses to the Economic Impact Payments of 2020: Evidence from the Consumer Expenditure Survey*. Working Paper 29648. National Bureau of Economic Research.
- Parker, Jonathan A. and Souleles, Nicholas S. (2019).** “Reported Effects versus Revealed-Preference Estimates: Evidence from the Propensity to Spend Tax Rebates.” *American Economic Review: Insights* 1.3, pp. 273–90.
- Pistaferri, Luigi (2001).** “Superior Information, Income Shocks, And The Permanent Income Hypothesis.” *The Review of Economics and Statistics* 83.3, pp. 465–476.
- Pratt, J. (1964).** “Risk Aversion in the Small and in the Large.” *Econometrica* 32.1/2, pp. 122–136.

- Rozsypal, Filip and Schlafmann, Kathrin (2023).** “Overpersistence Bias in Individual Income Expectations and Its Aggregate Implications.” *American Economic Journal: Macroeconomics* 15.4, pp. 331–71.
- Smith, Matthew, Zidar, Owen, and Zwick, Eric (Aug. 2022).** “Top Wealth in America: New Estimates Under Heterogeneous Returns.” *The Quarterly Journal of Economics* 138.1, pp. 515–573. ISSN: 0033-5533.
- Straub, L. (2019).** “Consumption, Savings, and the Distribution of Permanent Income.” *mimeo*.
- Toda, Alexis Akira (2021).** “Necessity of hyperbolic absolute risk aversion for the concavity of consumption functions.” *Journal of Mathematical Economics* 94.C.
- Wang, Chong, Wang, Neng, and Yang, Jinqiang (2016).** “Optimal consumption and savings with stochastic income and recursive utility.” *Journal of Economic Theory* 165, pp. 292–331. ISSN: 0022-0531.