

Heterogeneity in MPC Beyond Liquid Wealth: The Role of Permanent Earnings

Jeanne Commault*

May 2025

Abstract

While MPCs are mostly known to decrease with liquid wealth, I show that they are also increasing in the permanent component of earnings. In a standard model, permanent earnings raise MPCs because they reduce the ratio of risk-free-liquid-wealth-to-risky-future-earnings, strengthening precautionary behavior. This can explain two documented facts: (i) people with high levels of liquid wealth still have significant MPCs; (ii) MPCs do not decrease with current earnings although, like liquid wealth, they increase available resources. This prediction holds in survey data. The effect is large enough to explain the stylized facts. Numerical simulations match the survey results and stylized facts.

Key words: Marginal Propensity to Consume, Permanent Component of Earnings, Earnings Risk, Precautionary Saving, Standard Incomplete Market Model

JEL: D11, D12, D15, E21

1 Introduction

What determines people's marginal propensity to consume (MPC)? The MPC measures the fraction of a one-time, unexpected income change that is passed on to consumption over the period following the income change. Understanding how MPCs vary across people matters for several economic questions, including the transmission of economic shocks, the design of stabilization policies, and the dynamics of wealth accumulation.

An empirical result that emerged over the past two decades is that, on average, people respond significantly to a one-time unexpected shock (starting with studies including Parker 1999 and Johnson, Parker, and Souleles 2006). At first surprising, since consumers should smooth out the shock, several narratives have developed to explain this result. They tend to focus on the relation between

*Sciences Po, 28 Rue des Saints-Pères, 75007 Paris, France; jeanne.commault@sciencespo.fr.

liquid wealth and the MPC, since empirical studies have shown that liquid wealth significantly reduce MPCs. The seminal paper of Kaplan and Violante (2014) argues that, when distinguishing between liquid and illiquid wealth, the latter being costly to use for consumption smoothing, a large enough fraction of people hold either no wealth or only illiquid wealth to substantially impair consumption smoothing and obtain a realistically large average MPC. Other papers introduce different preferences (e.g. Aguiar, Bils, and Boar 2020) and behavioral biases such as temptation or present bias (e.g. Maxted, Laibson, and Moll 2024). These can also generate a realistically large average MPC, and their mechanism also typically involve a negative correlation between liquid wealth and the MPC: preferences or behavioral biases lead some consumers to have a lower desire to save and a higher MPC.

These models explain important aspects of the data but the more recent literature documents two stylized facts suggesting that additional factors are at play, beyond liquid wealth: (i) the extent to which liquid wealth reduces the MPC is significant but modest, so people with medium or high levels of liquid wealth have MPCs that are substantially above zero; (ii) higher current earnings do not associate with lower MPCs, although everything else equal people with higher current earnings have more immediate resources and could be less constrained in terms of liquidity; higher current earnings can even associate with higher MPCs.¹ These two facts hold for the response of total consumption to a shock but also for the response of nondurable consumption only.

In this paper, I show that a standard consumption model can account for those two stylized facts through the effect of the permanent component of earnings on the MPC. The permanent component of earnings captures the returns to permanent skills or abilities, and thus multiplies all realizations of an individual's earnings. I establish analytically, empirically, and numerically that, everything else equal, consumers with a higher level of permanent earnings have a higher MPC. This can explain stylized facts (i) and (ii): the positive effect of permanent earnings on the MPC partly offsets the negative effect of liquid wealth among consumers with more liquid wealth; the positive effect of the permanent component of current earnings on the MPC can dominate the negative effect of the more transitory components of current earnings so that, overall, current earnings do not reduce the MPC.

My first contribution is analytical. To highlight the mechanism at the core of the paper, I consider a simple version of the standard consumption model, in which a consumer solves an income fluctuation problem. Earnings evolve as a transitory-permanent process. I show that, in this model, at a given, positive, level of liquid wealth, an increase in the permanent component of earnings raises the MPC. Intuitively, the permanent component of earnings multiplies all current and future realizations of earnings. These future realizations are subject to shocks thus future earnings constitute a risky resource. Thus, at the same positive level of liquid wealth, consumers

¹Because there are many, I list the papers that establish those facts at the end of the introduction.

with a higher level of permanent earnings have more resources but a lower ratio of risk-free over risky resources. They consume more but save a larger fraction of their total expected resources because they seek to accumulate more risk-free wealth. Upon receiving a windfall that directly provides risk-free wealth, consumers with a higher permanent earnings, who would have saved more absent the shock, are more relieved and reduce their saving more. They still save more than others after a windfall, but the difference between their saving and that of others becomes smaller. Since their saving decreases more with a windfall, their consumption increases more: they have a higher MPC.

The effect of permanent earnings gets stronger when incorporating a consumption floor à la Stone-Geary. In that case, permanent earnings raise the MPC beyond their effect on the ratio of risk-free to risky resources: even comparing two consumers with the same ratio of wealth to permanent earnings, the one with higher permanent earnings has a higher MPC.

In the rarer case of consumers with negative liquid wealth, having a higher level of permanent earnings raises the ratio of liquid wealth to permanent earnings instead of reducing it. Thus, for people with negative liquid wealth, the MPC decreases with permanent earnings.

The generally positive effect of permanent earnings on the MPC can explain the stylized facts (i) and (ii). Consumers who have more liquid wealth are more likely to have a high permanent component of earnings. The positive effect of permanent earnings on the MPC then partly offsets the negative effect of liquid wealth. Absent a control for permanent earnings, the effect of liquid wealth appears smaller, explaining fact (i). Also, an increase in current earnings, which is made of a permanent component and of other more transitory components, may not reduce the MPC if the positive effect of the permanent component offsets the negative effect of the more transitory components, 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 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 into a permanent and a transitory innovation (Pistaferri 2001, Attanasio, Kovacs, and Molnar 2020), I show that they can also be used to identify the level of permanent earnings of the respondents. This is because expectations about future earnings are independent of the components of current earnings that are known to be transitory. I implement this method in the New York Fed Survey of Consumer Expectations (SCE). Consistent with the risk mechanism I highlight, the variance of future earnings increases with my measure of permanent earnings.

In a reduced form regression analysis, a one standard-deviation increase in permanent earnings

raises the reported quarterly MPC for total consumption by 0.09. Total consumption includes both durable and nondurable spending. This effect is statistically significant at the 5% level. It is half as large as the effect of a one-standard deviation decrease in liquid wealth. Thus, the magnitude is large enough that the effect of permanent earnings offsets a substantial fraction of the effect of liquid wealth, explaining fact (i). Consistent with this, when not controlling for permanent earnings, the effect of liquid wealth drops from 0.18 to 0.12 and is no longer precisely measured. Because the effect of permanent earnings is relatively large, it can also dominate the effect of the transitory component of earnings and generate fact (ii). Consistent with this, the coefficient associated with current earnings is positive but smaller than that of permanent earnings. Empirically, permanent earnings still raise the MPC when controlling for the wealth-to-permanent earnings ratio. This suggests the existence of a non-zero consumption floor à la Stone-Geary. One policy implication of this result is that targeting stimulus checks to people with lower income is not the most effective for a policy-maker seeking to maximize consumption response. Targeting based on a mix of wealth and age is more effective.

My third contribution is to show that numerical simulations of a standard incomplete market model can replicate my survey results and the stylized facts (i)-(ii) that motivate the analysis. Modeling earnings as proposed in Guvenen, Karahan, Ozkan, and Song (2021), with a realistic level of earnings risk, is key to obtain this.

I examine which components of the process in Guvenen, Karahan, Ozkan, and Song (2021) are the most important for the result. Having shocks drawn from mixtures of normal distributions rather than from normal distributions is sufficient to raise substantially the effect of permanent earnings on the MPC in the model.

Related literature. My results are most closely related to the empirical literature that examines the characteristics influencing people's MPCs. This is the literature that has established the stylized facts that constitute my starting point. What my paper brings to this literature is an explanation for these stylized facts, and an additional stylized fact that supports this explanation: that the permanent component of earnings has a significant and positive effect on the MPC, about half as large in absolute value as the effect of liquid wealth. More precisely, the stylized fact (i) that motivates my analysis and that has become a key finding in this literature is that, besides demographics, the main characteristic that consistently affects the MPC is liquid wealth, although the extent to which it reduces the MPC is modest enough that it does not move it down to zero. Some papers highlighting this include Jappelli and Pistaferri (2014), Baker (2018), Aydin (2019), Ganong, Jones, Noel, Farrell, Greig, and Wheat (2020), Fagereng, Holm, and Natvik (2021), Baugh, Ben-David, Park, and Parker (2021), or Golosov, Graber, Mogstad, and Novgorodsky (2023).² Recently, Bil-

²On this, Fuster, Kaplan, and Zafar (2020) write '... the only observable characteristic that has been robustly

biie, Galaasen, Gürkaynak, Mæhlum, and Molnar (2025) even document that the differences in MPC by level of liquid wealth (hand-to-mouth versus not) is much smaller when looking at MPCs out of disposable income. The stylized fact (ii) is that, in contrast to that of liquid wealth, the effect of total current earnings on the MPC is typically not significant. This is a result commonly documented in papers examining determinants of MPCs including Shapiro and Slemrod (2009), 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), and Boehm, Fize, and Jaravel (2025). Some papers even find a positive and significant effect of current earnings on the MPC. That is the case of Kueng (2018)³, Lewis, Melcangi, and Pilossoph (2022), and Kotsogiannis and Sakellaris (2025). 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 compatible with a positive correlation between current income and the MPC out of a shock of the same size for all. A recent working paper by Koşar and Melcangi (2025) presents new stylized facts on the relation between the MPC and the standard deviation of earnings growth. Here, I consider analytically and empirically the relation between the MPC and permanent earnings, through the fact that permanent earnings is proportional to the standard deviation of the earnings level—not of earnings growth.

My theoretical result builds on and extends the scope of the few analytical studies that have examined the role of the permanent component of earnings in the consumer's problem. My contribution is to examine its effect on the MPC and to introduce a departure from the homogeneous preferences case. Among them, Carroll (2006) notes that one can normalize the standard consumer's problem by the permanent component of earnings to solve the problem with one less variable. This is one of the results I use to prove that permanent earnings raises the MPC in the case with homothetic preferences. For ease of exposition of my proof, I expand the insight to show that the consumption solution of the problem is homogeneous of degree one in wealth and permanent earnings, not just scalable in permanent earnings.⁴ 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. Wang, Wang, and Yang (2016) proves the homogeneity of con-

shown to correlate with MPCs is holdings of liquid wealth, and even then the explanatory power of wealth for MPC heterogeneity is weak.', p1.

³The 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.

⁴The homogeneity is stronger than the normalization result: the normalization implies that the consumers' problem rewrites as a new one with the same properties but potentially different parameters when dividing all variables by permanent earnings; the homogeneity implies that the consumers' problem stays the same when dividing all variables by permanent earnings.

sumption in wealth and permanent earnings in a continuous model with Epstein-Zin preferences (p301). Straub (2019) proves the homogeneity of consumption in wealth and permanent income in a subcase of the standard consumption model with discrete time where agents are born with zero wealth and the permanent component of earnings is a time-invariant fixed effect (his Proposition 1). His paper then examines the effect of permanent earnings on the MPCP, that is, the concavity of consumption in permanent earnings.

Regarding my numerical results, the recent paper of Guvenen, Ozkan, and Madera (2024) shows that introducing the earnings process of Guvenen, Karahan, Ozkan, and Song (2021) has strong effects on the model's prediction, including an increase the average MPC generated by the model. I also find that introducing this process increases the MPC. Additionally, I show that it increases the impact of permanent earnings on the MPC. This stronger relation between permanent earnings and the MPC makes it possible for the model to generate the two stylized facts that motivate this paper: that the effect of liquid wealth is modest absent any control for permanent earnings, and that current earnings does not correlate negatively with the MPC.

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, I consider a simple version of the income fluctuations problem.

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 restrictions on the utility and earnings process, a budget constraint, and a terminal condition on wealth

$$V_t^i(a_t^i, e^{p_t^i}, e^{\varepsilon_t^i}) = \max_c u(c - c_0) + \beta E_t \left[V_{t+1}^i(a_{t+1}^i, e^{p_{t+1}^i}, e^{\varepsilon_{t+1}^i}) \right] \quad (2.1)$$

$$\text{with Isoelastic utility: } u'(\cdot) = (\cdot)^{-\rho}, \rho \geq 0 \quad (2.2)$$

$$\text{Earnings: } y_t^i = e^{p_t^i} e^{\varepsilon_t^i}, \quad \text{var}_{t-1}(\varepsilon_t) > 0 \quad (2.3)$$

$$\text{Permanent component: } e^{p_{t+1}^i} = e^{p_t^i} e^{\eta_{t+1}^i}, \quad \text{var}_{t-1}(\eta_t) \geq 0 \quad (2.4)$$

$$\text{Budget constraint: } a_{t+1}^i = (1 + r)a_t^i + y_t^i - c, \quad (2.5)$$

$$\text{Terminal wealth: } a_{T+1}^i \geq 0. \quad (2.6)$$

V denotes the value function. Utility is time-separable and, at each period t , depends on the difference between contemporaneous consumption expenditures c_t and a consumption floor c_0 . The period consumption utility function $u(\cdot)$ is isoelastic with $\rho \geq 0$ the relative risk aversion. It

thus displays constant relative risk aversion (CRRA). The discount factor β captures how much consumers discount utility between two consecutive periods.

The labor earnings specification, described with (2.3) and (2.4), is a transitory-permanent process: earnings are the product of a permanent component e^{p_i} that evolves as a multiplicative random walk and of a transitory innovation e^{ε_i} that is an i.i.d. shock. Because e^{p_i} , the permanent component at t , multiplies the value of the permanent component at $t + 1$, which itself multiplies the permanent component at $t + 2$, it multiplies each future realization of earnings until the rest of the consumer's lifetime. For each $t + s$

$$y_{t+s}^i = e^{p_i} e^{\eta_{t+1}^i + \dots + \eta_{t+s}^i} e^{\varepsilon_t^i}.$$

The permanent component of earnings e^{p_i} thus scales all future realizations of earnings. Note that this specification encompasses an even simpler specification in which the permanent component is just a multiplicative fixed effect $e^{p_i} = e^{p^i}$. Throughout this paper, it is not necessary that consumers face permanent income shocks e^η . I assume that they face some uncertainty about their future income so that they have a precautionary saving motive. However, uncertainty coming only from future transitory shocks is sufficient: I only impose $\text{var}(\varepsilon) > 0$ and let $\text{var}(\eta) \geq 0$. 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.

There is only one asset, that is risk-free and perfectly liquid. The budget constraint (2.5) describes the evolution of this risk-free liquid wealth from a period to the next. The term a_t^i denotes the level of risk-free liquid wealth at the beginning of period t —or at the end of $t - 1$. The risk-free return rate is r .

The terminal condition on wealth (2.6) 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 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 an isoelastic utility implies that 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.

Euler equation. The intertemporal first-order condition of the optimization problem described by (2.1)-(2.6), known as the Euler equation, states that

$$u'(c_t^i - c_0) = E_t[u'(c_{t+1}^i - c_0)]R, \quad (2.7)$$

where $R \equiv \beta(1+r)$ is the factor that account for the difference between the individual's discount factor β and that of the market $1/(1+r)$.

MPC definition. I want to examine the effect of permanent earnings on the MPC, where the MPC is the response of consumption to an unexpected one-time income shock. Such a shock would be modeled as a term w_t unexpectedly 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 exactly the same effect on the consumption decision as an unexpected change in beginning-of-period wealth a_t of the same magnitude: in this model an unexpected shock w_t —such as a stimulus check—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

Proposition 1. In the model described above by (2.1)-(2.6), when $R \leq 1$, at any period $t < T$, an increase in the permanent component of earnings e^{p_t} strictly raises the MPC of consumers with strictly positive risk-free liquid wealth $a_t > 0$

$$\frac{\partial MPC_t}{\partial e^{p_t}} = \frac{\partial^2 c_t}{\partial a_t \partial e^{p_t}} > 0 \text{ when } a_t > 0.$$

In contrast, it strictly reduces the MPC of consumers with strictly negative liquid wealth.

Proof of Proposition 1. In the case without any consumption floor ($c_0 = 0$), the result builds on the combination of two existing results. First, with isoelastic utility and no consumption floor, Carroll (2006) notes that the consumer problem scales in permanent earnings, that is, it can be rewritten as a similar looking problem with all variables divided by the level of permanent earnings. Although the scalability result is sufficient to obtain Proposition 1, I prove for convenience of the reasoning that the consumption solution of the consumer's problem is not just scalable in permanent earnings but homogeneous of degree one in wealth and permanent earnings (Lemma 4, Appendix ??).⁵ I

⁵Scalability is not the same as homogeneity of degree one. Under scalability, if $c_t = f^t(a_t, e^{p_t}, e^{e_t})$, then there exists

denote $c_t = f^t(a_t, e^{p_t}, e^{\varepsilon_t})$ the consumption function. Its homogeneity of degree one in a_t and e^{p_t} implies $c_t/e^{p_t} = f^t(a_t/e^{p_t}, 1, e^{\varepsilon_t})$. Differentiating $c_t = e^{p_t} f^t(a_t/e^{p_t}, 1, e^{\varepsilon_t})$ with respect to a_t I have

$$MPC_t = \frac{\partial c_t}{\partial a_t} = f_1^t\left(\frac{a_t}{e^{p_t}}, 1, e^{\varepsilon_t}\right), \quad (2.8)$$

where $f_1(\cdot, \cdot, \cdot)$ is the partial derivative of the consumption function with respect to its first argument. The MPC is homogeneous of degree zero in a_t and e^{p_t} : if one divides both a_t and e^{p_t} by the same non-zero value, the MPC remains unchanged.

Second, Carroll and Kimball (1996) proves that when the utility function displays Hyperbolic Absolute Risk Aversion (HARA)—a set that includes all isoelastic functions—and $c_0 = 0$, consumption is strictly concave in liquid wealth. This means that the MPC is decreasing in wealth. For any a_t , e^{p_t} , and e^{ε_t}

$$\frac{\partial MPC_t}{\partial a_t} = \frac{\partial^2 c_t}{\partial a_t^2} = \frac{1}{e^{p_t}} f_{11}^t(a_t, e^{p_t}, e^{\varepsilon_t}) < 0. \quad (2.9)$$

where $f_{11}(\cdot, \cdot, \cdot)$ is the second-order partial derivative of the consumption function with respect to its first argument. I show in Commault (2025) that, when utility is not just HARA but isoelastic (i.e. CRRA), the reason why the MPC is lower at a higher level of wealth is because an increase in wealth reduces precautionary saving but less so at a higher level of wealth. With a HARA but non-CRRA utility function, the consumption function is concave in wealth but an increase in wealth does not necessarily reduce precautionary saving.

Combining these two results implies that the MPC is increasing in permanent earnings e^{p_t} when $a_t > 0$. To show this, I differentiate both sides of (2.8) with respect to e^{p_t}

$$\frac{\partial MPC_t}{\partial e^{p_t}} = \frac{\partial^2 c_t}{\partial a_t \partial e^{p_t}} = -\frac{a_t}{(e^{p_t})^2} f_{11}^t\left(\frac{a_t}{e^{p_t}}, 1, e^{\varepsilon_t}\right) = -\frac{a_t}{e^{p_t}} \frac{\partial MPC_t}{\partial a_t} > 0. \quad (2.10)$$

Because the MPC only depends on the ratio e^{p_t}/a_t , the effect of e^{p_t} on the MPC goes in the opposite direction as the effect of a_t when $a_t > 0$. Because the MPC decreases with a_t , then it must increase with e^{p_t} . Intuitively, experiencing an increase in permanent earnings is equivalent to experiencing a proportional increase in both wealth and permanent earnings (which keeps the MPC constant when it is homogeneous) and a decrease in wealth back to its initial level (which raises the MPC). Thus, an increase in permanent earnings raises the MPC. Since the reason why the MPC increases with a decrease in wealth is because a decrease in wealth raises precautionary saving and makes this saving more sensitive to windfalls, the reason why the MPC increases with an increase

a policy function $\tilde{f}^t$ with the same properties as f^t but possibly different parameters such that $c_t/e^{p_t} = \tilde{f}^t(a_t/e^{p_t}, 1, e^{\varepsilon_t})$. Under homogeneity of degree one, the two policy functions coincides with $c_t/e^{p_t} = f^t(a_t/e^{p_t}, 1, e^{\varepsilon_t})$.




Figure 1: The evolution of consumption with wealth at two levels of permanent earnings

in permanent earnings is the same: an increase in permanent earnings raises precautionary saving and makes this saving more sensitive to windfalls.

The opposite is true when $a_t < 0$: when people have debt, that is, when net wealth is negative, an increase in permanent earnings decomposes into a proportional increase in both permanent earnings and debt (which keeps the MPC constant), and a decrease in debt thus an increase in wealth (which reduces the MPC). In that case, having a higher level of permanent earnings has the same impact on the MPC as having less debt. It reduces the MPC: $\partial MPC_t / \partial e^{P_t} = -(a_t / (e^{P_t})^2) \tilde{c}_2'(a_t / e^{P_t}, 1, e^{E_t}) < 0$.

In the more general case with a non-zero consumption floor $c_0 \geq 0$, the reasoning presented above still applies. Indeed, it is not necessary that the MPC be homogeneous of degree zero and remain unchanged when both permanent earnings and liquid wealth increase proportionally. It is sufficient that it either remains unchanged or increases. I prove in Lemma 6, Appendix ?? that, in the case with $c_0 \geq 0$, the MPC increases when both $a_t > 0$ and e^{P_t} increase proportionally. I also prove in Lemma 7, Appendix ?? that, in the case with $c_0 \geq 0$, c_t is still strictly concave in a_t . An increase in permanent earnings e^{P_t} is equivalent to a proportional increase in a_t and e^{P_t} (which either keeps the MPC constant when $c_0 = 0$ or raises it when $c_0 > 0$) and a decrease in a_t back to its initial level (which raises the MPC). Thus, in the case with $c_0 \geq 0$, an increase in permanent earnings still raises the MPC. Furthermore, when $c_0 > 0$, an increase in permanent earnings raises the MPC even if wealth a_t increases proportionally with permanent earnings. The section titled Proposition 1 in Appendix ?? details how the combination of the two Lemmas establishes Proposition 1.

Graphical illustration of Proposition 1. Figure 1 presents graphically the mechanism. It plots the evolution of consumption with wealth conditional on all other determinants of consumption. The different lines correspond to different levels of permanent earnings (baseline in black versus high in blue) and different levels of uncertainty about future earnings (no uncertainty in dashed and some uncertainty in plain). The slope of each line corresponds to the MPC since it measures by how much consumption increases when wealth increases by one unit.

The dashed black line represents the evolution of consumption with wealth for a baseline level of permanent earnings and under perfect foresight, that is, absent any uncertainty about future earnings. The evolution of consumption is a straight line: under perfect foresight the optimum is to consume a fixed fraction of lifetime expected resources—for instance this fraction would be about 3% for consumers with 60 periods left to live, a discount factor of 0.98, and facing an interest rate of 2%. One extra unit of wealth raises consumption because it raises lifetime resources, and it raises it by the same amount at each level of wealth. The MPC is constant across the wealth distribution and equal to the fixed fraction of their resources that the consumer consumes. This perfect foresight case is also the 'permanent income' case considered in Friedman (1957)—where permanent income denotes the amount that people consume, not the permanent component of a permanent-transitory earnings specification. This consumption function also coincides with the one of an individual with quadratic preferences (Hall 1978).

The plain black line plots the evolution of consumption with wealth for the same baseline permanent earnings as in the dashed line, but consumers now face uncertainty from earnings shocks. The difference between the dashed and plain lines corresponds to precautionary saving: it measures the extra saving caused by earnings uncertainty. The slope is steeper than under perfect foresight because an increase in wealth now raises consumption for two reasons: it raises lifetime expected resources, as in the dashed case, and it also reduces the need for precautionary saving. The MPC is thus higher than under perfect foresight and it varies along the wealth distribution: one additional unit of wealth relaxes the need to make precautionary saving more at a lower level of wealth, so the MPC is larger at a lower level of wealth. The slope of the consumption function decreases with wealth, making consumption concave.

I now plot in blue a situation in which permanent earnings e^{pt} is higher than in the baseline (black) situation. The other determinants of consumption, including the transitory component ε_t and distribution of future shocks ε_{t+s} and η_{t+s} , are the same as in the black situation. At each level of wealth a_t , expected lifetime resources are now larger. Under perfect foresight (dashed blue line), consumption is simply shifted up: the consumer consumes the same fraction of a larger level of lifetime resources. The dashed blue line is strictly above the dashed black line but they have the same slope.

In the presence of earnings uncertainty (plain blue line), the consumer also consumes more at a

higher level of permanent earnings than at the baseline level: the plain blue line lies above the plain black line. However, the slopes are different. Because permanent earnings raises precautionary saving and makes it more sensitive to variations in wealth, the precautionary gap between the plain and dashed line is more pronounced in the blue situation, and one additional unit of wealth reduces precautionary saving more in the blue situation. As a result, at the same level of wealth, the slope of the plain blue line is steeper than the slope of the plain black line: the MPC is higher at a higher level of permanent earnings.

Note that, in the case without a consumption floor, the MPC is homogeneous of degree zero in wealth and permanent earnings. This means that increasing permanent earnings and wealth proportionally (moving from the black to the blue line and also moving right on the x-axis by an amount that would be proportional the underlying increase in permanent earnings) would keep the slope of the line unchanged. In the case with a non-zero consumption floor, the slope of the plain blue line would still be steeper than the slope of the plain black line after increasing both permanent earnings and wealth proportionally.

When is it sufficient to substitute wealth with the ratio of wealth to permanent earnings in MPC analyses? In both empirical studies and numerical simulations, people focus on liquid wealth as one of the most important determinant of the MPC apart from demographic characteristics. Proposition 1 shows that permanent earnings is another determinant. In the case without any consumption floor, to take Proposition 1 into account, it is sufficient to substitute liquid wealth with the ratio of liquid wealth to permanent earnings in empirical analyses. In the case with a non-zero consumption floor, the ratio is no longer a sufficient statistics. Changes in permanent earnings and wealth affect the MPC even when their ratio is unchanged. Taking Proposition 1 into account then requires considering the effects of both liquid wealth and permanent earnings in empirical analyses.

Discussion of exogenous borrowing constraints. This theoretical section relies on a simple model to highlight the positive relation between permanent earnings and the MPC that the simplest version of the consumer problem generates, without exogenous borrowing constraints. The mechanism works through precautionary saving.

Under some conditions, Proposition 1 can extend to the presence of exogenous liquidity constraints. That is the case when the exogenous borrowing constraint is proportional to the permanent component of earnings—for instance because the bank lets consumer with a higher level of permanent earnings borrow more—. In that case, the consumer’s problem remains homogeneous of degree one. In addition, I know from Carroll, Holm, and Kimball (2021) that consumption remains concave in wealth in the presence of uncertainty and borrowing constraints. As a result, an increase in permanent earnings, akin to a proportional increase in both permanent earnings

and wealth (keeping the MPC constant from the MPC homogeneity) and a decreasing in wealth (reducing the MPC from the concavity of consumption in wealth) still reduces the MPC.

When the exogenous borrowing constraint is fixed rather than proportional to permanent earnings, its effect on the MPC depends on whether an increase in permanent earnings strengthens the borrowing constraint or relaxes it. Which scenario prevails depends on whether earnings are sufficiently increasing over the life-cycle—in which case an increase in permanent earnings may increase future resource enough to make people want to move resources from the future to the present more and strengthen the constraint—or remain flat enough over the life-cycle—in which case the fact that an increase in permanent earnings raises current cash-in-hand may prevail and relax the constraint.

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 (2015-2019). It is a monthly online survey with a rotating panel of about 1,300 household heads based in the United States. 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. The Labor Survey module, which reports earnings information, takes place every four months. The Housing module and the Household Finance survey, which contain information about wealth, take place once a year.⁶ In addition to these topical modules Fuster, Kaplan, and Zafar (2020) fielded MPC-specific modules at four points in time between March 2016 and March 2017, which I use to measure the respondents' MPCs. My period of observation is thus between March 2016 and March 2017. Half of the MPC-specific modules, and more than half of the observations, take place on the same month as a Labor Survey module. When that is not the case, I match the MPC reported in the MPC modules with the closest previous observations of earnings variables. I detail the way in which I match the modules in Appendix ??.

Earnings and future employment probability. I obtain current annual earnings, expected future annual earnings, and the probability to be employed in the future from questions in the Labor Mar-

⁶See Armantier, Topa, Klaauw, and Zafar (2017) for technical background information on the SCE, and www.newyorkfed.org/microeconomics/sce.html for additional information.

ket 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 non-employed. This makes it possible to build a measure of the variance of future earnings as foreseen by the individuals themselves, which I use to verify that the standard-deviation of future earnings increases linearly with permanent earnings in my sample of employed respondents.

MPCs. To measure MPCs, I use the modules added by Fuster, Kaplan, and Zafar (2020) to the end of the monthly surveys in March 2016, May 2016, January 2017 and March 2017. Respondents are asked to report how they would change their spending behavior in response to an unexpected change in resources. The baseline question is about a \$500 gain today. The hypothetical situation is framed as follows

Now consider a hypothetical situation where you unexpectedly receive a one-time payment of \$500 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.

People are then asked about their response in two steps. In a first step, they describe qualitatively whether they would change anything over the next 3 months to their spending and donation, to their level of debt, and to their saving. In a second step, people are asked to quantify the changes they would make. This two steps procedure generally reduces the extent to which people select answer only at 0, 0.5 and 1 because the first step leads them away from corner solutions. Besides the \$500 gain today, people are asked about their spending response to a \$2500 gain today, a \$5000 gain today, a \$500 gain in 3 months, a \$5000 gain in 3 months, a \$500 loss today, a \$500 loss in 3 months, and a \$500 loss in 2 years. The MPC corresponds to the reported change in spending divided by the size of the stated gain or loss. There is no constraint that the MPC be smaller than one nor larger than zero. 4% of observations correspond to a MPC strictly below zero or strictly above one after trimming the top and bottom 1% of the distribution (see the last two lines of Table ?? in Appendix ??). In the analysis, I pool together MPCs from questions about different types of shocks and use dummies to control for the effect of the MPC question, with a \$500 gain today as the reference category. On this, the model prediction is that the response to all these questions should be larger at a higher level of permanent earnings, whether it is a positive or a negative shock, whether it is a big or small shock, and whether it is a shock today or in the future. I still examine the heterogeneity by type of MPC question in robustness analyses. The only question in the MPC modules that I exclude is about the change in spending in response to a \$5000 zero-interest loan today to be repaid one year from now. The same respondent can be asked several of these questions in the same module. An observation is the response of a given respondent to a given MPC question at a given date.

The resulting average reported MPC is 0.176. This MPC includes all spending over the three months following the shock. This is in line with, though on the lower end of, recent results measuring MPCs from natural experiments, where the natural experiment is typically a staggered disbursement of tax rebates. While a first range of studies obtained point estimates for the quarterly MPC of total consumption above 0.50 (see e.g. Johnson, Parker, and Souleles (2006) Parker, Souleles, Johnson, and McClelland (2013), Broda and Parker (2014)), more recent studies using econometric techniques robust to treatment effect heterogeneity, put the quarterly MPC of total consumption below. Orchard, Ramey, and Wieland (2025) and Borusyak, Jaravel, and Spiess (2024) find the quarterly MPC for total consumption out of the 2008 US Economic Stimulus Payments to be around 0.25. Parker, Schild, Erhard, and Johnson (2022) estimate the quarterly MPC out of the covid stimulus payment to be 0.10 for nondurable spending, with little evidence that durable responded. In an experimental design where participants receive a 300 euros cash transfer, Boehm, Fize, and Jaravel (2025) estimate the MPC of total spending out of the transfer to be 0.23 after one month. The MPCs reported in the modules of Fuster, Kaplan, and Zafar (2020) are also in line with result from other surveys asking about the response to an hypothetical shocks (e.g. Graziani, Klaauw, and Zafar (2016) Christelis, Georgarakos, Jappelli, Pistaferri, and Rooij (2019), Bunn, Le Roux, Reinold, and Surico (2018), Parker and Souleles (2019) Crossley, Fisher, Levell, and Low (2021)).

Parker and Souleles (2019) and Kotsogiannis and Sakellaris (2025) find that the MPCs reported in hypothetical scenarios is consistent with what people report consuming out of a realized fiscal stimulus and with what people do upon receiving a one-time lottery winning respectively. Colarieti, Mei, and Stantcheva (2024) validate the reliability of surveys in predicting actual economic behaviors using a cross-validation method.

I also compute the yearly MPC. I use questions in the Fuster, Kaplan, and Zafar (2020)'s modules about the amount spent on each of the month of the quarter following the shock to derive the rate at which the monthly spending declines over time. I then apply this rate to the following nine months and obtain a yearly MPC of 0.477.

Liquid wealth. I use two measures of liquid wealth. The first one is from a question in the Housing module of the SCE asking respondents to select which category of net non-housing wealth their household belongs. They are offered fourteen possible bins, ranging from below five hundred dollars to above one million dollars. The advantage of this question is that it takes place around the same period as the MPC modules and a large number of respondents of these modules answer it. The second measure is from two questions in the Household Finance module: one about the amount of savings and investments in accounts other than retirement accounts, and one about the share of this amount hold in saving and checking accounts. This measure more precisely captures

liquid wealth. However, the survey takes place in August, further away on average from the dates MPC modules, so the number of MPC survey respondents who answer it is lower.

Demographics. The demographics, including age, gender, education level, willingness to take risks, number of household members, number of children in the household, and household income categories, are 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 assume away the risk of non-employment since employed respondents still 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 the survey date and characteristics of the respondent. This does not seem to hold when I include non-employed respondents, who appear to draw earnings shocks from riskier distributions than employed respondents. 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 above age 55. Finally, I winsorize the top and bottom 1% of expected future earnings, earnings, the MPC, consumption, and the variance of future earnings variables. I obtain a final sample of 2,733 observations where I observe jointly the main variables. These observations are coming from 905 respondents. Analyses are then clustered at the respondent level. I deflate all the \$ variables using the monthly non-seasonally adjusted Consumer Price Index (CPI) to express them all in March 2016\$.

Tables. I present the text of the questions and detail the way I build my main variables in Appendix ???. I present descriptive statistics of these variables in this final sample in Tables ?? and ??, Appendix ??.

3.2 Measuring permanent earnings

Identification insight. Consider a general earnings specification such that annual earnings of consumer i at $t + 1$ is the product of a permanent component $e^{p_{t+1}^i} = e^{p_t^i} e^{\eta_{t+1}^i}$ that evolves as a multiplicative random walk and of a term x_{t+1}^i . Expected annual earnings is

$$y_{t+1}^i = e^{p_t^i} e^{\eta_{t+1}^i} x_{t+1}^i \Rightarrow E_t^i[y_{t+1}^i] = e^{p_t^i} E_t^i[e^{\eta_{t+1}^i} x_{t+1}^i] \quad (3.1)$$

Assuming the term $E_t^i[e^{\eta_{t+1}^i} x_{t+1}^i]$ is the same for respondents with the same observed characteristics at t , I can filter out variations in this term with demographic and date controls. Differences in expected future earnings $E_t^i[y_{t+1}^i]$ that are not explained by differences in demographics and date

are then driven by differences in permanent earnings e^{p_i} . That is the intuition for the identification.

A generalized transitory-permanent earnings process. I show that this insight applies to one of the generalized transitory-permanent earnings process proposed in Guvenen, Karahan, Ozkan, and Song (2021)—their specification (5). This specification encompasses the simple transitory-permanent process that is common in the literature and that I use in section 2 to illustrate the mechanism. Guvenen, Karahan, Ozkan, and Song (2021) find that their specification fits well the moments of annual earnings measured from administrative US data

$$\text{Annual earnings: } y_t^i = \underbrace{(1 - v_t^i)}_{\text{Employment status}} \underbrace{e^{\tilde{\alpha}^i + \zeta z^i}}_{\text{Fixed effect}} \underbrace{e^{p_t^i}}_{\text{Highly persistent}} \underbrace{e^{\varepsilon_t^i}}_{\text{Transitory}} \underbrace{e^{g(t)}}_{\text{Age trend}} \quad (3.2)$$

$$\text{Persistent component: } e^{p_t^i} = (e^{p_{t-1}^i})^\rho e^{\eta_t^i}, \quad (3.3)$$

$$\text{Nonemployment: } v_t^i \sim \begin{cases} 0 \text{ (employment) with prob. } 1 - p_{v_{t-1}^i}, \\ 1 \text{ (nonempl.) with prob. } p_{v_{t-1}^i}. \end{cases} \quad (3.4)$$

This expression states that the annual earnings of worker i at t , y_t^i , 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 can include a part depending on observed fixed individual component 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 and ρ its persistence. In practice, this component is virtually permanent because Guvenen, Karahan, Ozkan, and Song (2021) estimate its persistence to be $\rho = 0.991$.⁷

Defining permanent earnings. My objective is to capture the part of this earnings process that is 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 $e^{p_t^i}$ 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 to ease the interpretation of its unit change: I re-scale it so that a 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 earnings of respondent i at period t

$$perm_t^i = e^{\tilde{\alpha}^i + \zeta \bar{z}} e^{p_t^i} e^{\bar{\varepsilon}} e^{\bar{g}} \quad (3.5)$$

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

⁷Furthermore, Rozsypal and Schlafmann (2023) find that people are partly overestimating the persistence ρ of their earnings. Since the ρ that I should use is the one that people believe rather than the realized one, it is likely to be even higher than the realized one.

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

$$\underbrace{\ln\left(\frac{E_t^i[y_{t+1}^i]}{(1-p_{v_t}^i)}\right)}_{\text{Observed}} = \underbrace{\overbrace{\rho p_t^i + \tilde{\alpha}^i}^{res_t^i}}_{\approx 1} + \underbrace{\zeta z^i + \ln(E_t^i[e^{\eta_{t+1}^i}]) + \ln(E_t^i[e^{\varepsilon_{t+1}^i}]) + g(t+1)}_{\text{Captured through fixed individual characteristics and year dummies}}. \quad (3.6)$$

If individual characteristics 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 $\ln(E_t^i[e^{\varepsilon_{t+1}^i}])$ across respondents are captured by a linear regression over such dummies.⁸ As a result, the residual res_t^i from a regression of $\ln(E_t^i[y_{t+1}^i]/(1-p_{v_t}^i))$ on individual characteristics and year dummies coincides with $p_t^i + \tilde{\alpha}^i$.⁹ In the baseline, the individual characteristics z^i that I control for are the gender of the respondent and their willingness to take risks. Studies document that both are likely to affect the distributions from which people draw shocks.

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, denoted $\overline{\ln(y)}|_{v_t^i=0} = \bar{\varepsilon} + \zeta \bar{z} + \bar{g}$. I then take the exponential and I obtain what I aim to measure: the permanent component of earnings

$$e^{res_t^i \times \overline{\ln(y)}|_{v_t^i=0}} = e^{p_t^i + \tilde{\alpha}^i} e^{\bar{\varepsilon}} e^{\delta \bar{z}} e^{\bar{g}} = perm_t^i.$$

In the survey data, I observe earnings over a longer period than the MPCs. To gain precision in my measure of permanent earnings I build it with March 2015-March 2019 data. I then only use observations of permanent earnings between March 2016 and March 2017. Table ?? in Appendix ?? 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 but to identify the permanent component of the level of earnings rather than the permanent component of the shocks. I note that when ex-

⁸This is because the log of the expected values $\ln(E_t^i[e^{\eta_{t+1}^i}])$ and $\ln(E_t^i[e^{\varepsilon_{t+1}^i}])$ 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 $\ln(E_t^i[e^{\varepsilon_{t+1}^i}]) \approx E_t^i[\varepsilon_{t+1}^i] + \text{Var}_t^i(\varepsilon_{t+1}^i)/2$.

⁹Indeed, it coincides with $p_t^i - \bar{p} + \tilde{\alpha}^i - \bar{\alpha}$ and 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 2: The Standard Deviation of Future Earning And Permanent Earnings

pected future employment is observed, my method generalizes to a process with non-employment shocks. Other methods that rely on subjective expectation data include Arellano, Attanasio, Crossman, and Sancibrián (2024) and Arellano, Attanasio, De Nardi, Borella, and Paz-Pardo (2024). These methods aim at estimating the parameters of earnings processes. Here, I do not estimate the parameters of a process, e.g. the persistence or the cumulative distribution function of the shock distribution of permanent earnings. What I identify is the value of permanent earnings for each respondent. Other methods to proxy the permanent component of earnings include simply using current earnings or using an average of past and current earnings. A method which, like mine, allows for a general earnings specification with non-employment shocks is that of Braxton, Herkenhoff, Rothbaum, and Schmidt (2021) who develop a filtering algorithm. To work well, this requires a longer time series than I have in the SCE.

The effect of heterogeneous trends. In the baseline, I draw from specification (5), the simpler of the two specifications that Guvenen, Karahan, Ozkan, and Song (2021) identify as fitting the data well. Specification (6) additionally includes an heterogeneous profile $e^{\beta^i t}$ in the realization of annual earnings at t . The term β^i is distributed with mean zero. In the presence of such profiles, what I measure with my method is $pe\hat{m}_t^i = e^{\tilde{\alpha}^i + \zeta \bar{z} + \beta^i(t+1)} e^{p_t^i} e^{\bar{e}} e^{\bar{g}} = perm_t^i e^{\beta^i(t+1)}$. This term also multiplies all future realizations of earnings, although a future realization at $t + s$ is not multiplied only by $pe\hat{m}_t^i$ but also by $e^{\beta^i(s-t-1)}$. If the profile is increasing over time, so that $e^{\beta^i(s-t-1)}$ is large, the permanent component of earnings at t has an even bigger multiplicative effect on future earnings than on current earnings. This should strengthen my effect for this group. Overall, the effect of permanent earnings on the MPC should then be stronger for a subgroup with steeper profiles and lower for a subgroup with profiles less steep than others.

Permanent earnings and the standard deviation of future earnings. As a first check of my theo-

retical mechanism, I examine how the standard deviation of future earnings varies with permanent earnings. One typical difficulty with measuring risk is that the ex-ante distribution of outcomes that an individual faces before the risk realizes is not observed. Only the one outcome that realizes ex-post is. A typical solution is to proxy the individual earnings risk of a given respondent with the variance of ex-post realized earnings innovation among a group of similar respondents. The SCE makes it possible to solve this problem without making this assumption and obtain directly an individual-level measure of the variance of future earnings. This is because it asks respondents about the probabilities that they assign to different events that would change their annual earnings. The survey questions that I use are about

- the probability to receive job offers and the annual earnings associated with the best offer
- the probability to accept these job offers
- the probability that their current employer matches these job offers
- the probability to become unemployed or to leave the labor market.

From these I build, for each worker, the set of possible future annual earnings they may have at the next period, and the probability they attach to each level of future annual earnings. I detail the questions and the way I combine them in Appendix (??). I then build the variance of future annual earnings of a given worker as the sum of this worker's possible levels of future annual earnings squared, weighted by their probability. The standard deviation of permanent earnings is the square root of this variance. Note that the questions I use to build the standard deviation of future earnings do not include the expected future annual earnings question that I use to build permanent earnings so no relation between the variables is built in.

Once I have the permanent earnings and variance of future earnings of each respondent, I discretize permanent earnings in fifteen 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 2 plots the relation between the two. It shows that people's standard deviation of future earnings increases linearly with their permanent earnings. This is what an earnings process with a permanent component predicts: if $y_{t+1} = perm_t x_{t+1}$ then $var_t(y_{t+1}) = perm_t^2 var_t(x_{t+1})$ and $sd_t(y_{t+1}) = perm_t sd_t(x_{t+1})$. Incidentally, the earnings specification of Guvenen, Karahan, Ozkan, and Song (2021), which I show is consistent with my permanent earnings identification, only implies a proportional and positive relation between permanent earnings and the standard deviation of future earnings *conditional on being employed in the future*.¹⁰ Figure 2 shows that the broadly

¹⁰Formally, the relation is $sd_t(y_t^i) = perm_t^i \sqrt{(1 - p_{vi}^i)} e^{\delta(z^i - \bar{z})} e^{g(t+1) - \bar{g}} sd_t(e^{\eta_{t+1}^i} e^{\varepsilon_{t+1}^i - \bar{\varepsilon}})$. When a high permanent earnings protects from unemployment, the positive correlation between $\sqrt{(1 - p_{vi}^i)}$ and $perm_t^i$ may change the relation between $sd_t(y_t^i)$ and $perm_t^i$ away from a linear relation.

linear relation holds even between permanent earnings and the *unconditional* standard deviation of future earnings. This means that, although I find that a higher level permanent earnings reduces reported future unemployment risk in the data, this effect is not large enough to strongly alter the linear relation between permanent earnings and the standard deviation of future earnings.

Permanent earnings and the coefficient of variation of future earnings. In addition, when earnings evolve as a transitory-permanent earnings process, the theoretical prediction is that there is no relation between the coefficient of variation of future earnings, defined as the ratio of the standard deviation over the mean, and current earnings. To see this, under the specification of Guvenen, Karahan, Ozkan, and Song (2021), the coefficient of variation is $CV_t^i = sd_t^i(e^{\eta_{t+1}^i} e^{\varepsilon_{t+1}^i - \bar{\varepsilon}}) / E_t^i[e^{\eta_{t+1}^i} e^{\varepsilon_{t+1}^i - \bar{\varepsilon}}]$. It depends on the expected distribution of future shocks but not on y_t^i . I find that, as predicted, the coefficient of variation does not vary with current earnings in my baseline sample where I exclude the non-employed. However, it decreases with earnings when I include both non-employed and employed respondents. I describe those results this in Appendix ???. This is presumably because non-employed respondents, with close to zero earnings, draw shocks from riskier distributions than employed respondents. As a result, the coefficient of variation decreases when earnings rise away from zero. This latter result is consistent with the findings in Arellano, Bonhomme, Vera, Hospido, and Wei (2021), who find that the coefficient of variation of future earnings decreases with current earnings, but that the decrease is essentially driven by people with low attachment to the labor market.

Ruling out anticipations and verifying that expected annual earnings differs from current earnings 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 $perm_t^i \times e^{\varepsilon_{t+1} - \bar{\varepsilon}}$. To 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 ε_{t+1} is anticipated, it will present both in my measure of permanent earnings and in the realized innovation to log-earnings: the two would covary. Contrary to that, I find that their correlation is quantitatively very small and not significant. I present those results in Appendix ???. Note that if people anticipate correctly the value of their permanent shock, this shock will also enter my measure of current permanent earnings. This is not a problem since this shock corresponds to permanent earnings and has the same impact on the MPC as realized permanent earnings.

Another problem could be respondents reporting the same value for their current earnings and for their expected future earnings. I examine whether that is the case. I find that people do not report the same values for current annual earnings and expected future annual earnings (see Table

?? in Appendix ??). The horizon of the question seems long enough for respondents to expect some change.

4 Permanent earnings and the MPC in survey data

4.1 Specification and main results

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

$$MPC_t^i = a_1 + a_2 \text{perm cat}_t^i + a_3 \text{nonhousing cat}_t^i + a_4 X_t^i + \xi_t^i \quad (4.1)$$

$$MPC_t^i = b_1 + b_2 \text{perm}_t^i + b_3 (\text{perm}_t^i)^2 + b_4 \text{checking}_t^i + b_5 (\text{checking}_t^i)^2 + b_6 X_t^i + \chi_t^i \quad (4.2)$$

The term MPC_t^i is the reported MPC out of a hypothetical shock of respondent i at period $\times$ MPC question t .¹¹ The term perm cat_t^i is a vector of dummies for the quartile of the permanent earnings distribution, where permanent earnings is built as described above. The reference category is the first quartile of permanent earnings. The term $\text{nonhousing cat}_t^i$ is a vector of dummies for quartiles of the non-housing household wealth distribution. These quartiles are obtained from merging a categorical measure of non-housing household wealth into four new categories that broadly correspond to the quartile of the distribution in my sample. The reference category is the first quartile of non-housing wealth.¹² The term X_t^i designates a vector of control variables. I show results for different sets of controls. The baseline controls are dummies for the MPC question that the respondent is asked, interacted with the month-year in which the question is asked. This corresponds to thirteen dummies. The reference category is a \$500 gain today, asked in March 2016.¹³ The respondent controls are dummies for the age group (eight categories from 'below 25' to '50-55'), the gender, and the education level (high-school, some college, completed college) of the respondent. The household controls include the number of household members, the number of children (below 18) in the household, and the categories household pre-tax income of past year (eleven categories from 'Less than \$10,000' to '\$200,000 or more'). I no longer use household controls for those from specification (4) on, in which the sample is of single-person households

¹¹ An observation is a respondent $\times$ period $\times$ MPC question but instead of keeping an extra letter to index the MPC question I use t to denote the period $\times$ MPC question.

¹² The first quartile corresponds to 'Less than \$5,000'. The second quartile is 'Between \$5,000 and \$20,000'. The third quartile is 'Between \$20,000 and \$50,000'. The fourth quartile is 'More than \$50,000'.

¹³ The list of dummies is as follows: \$5000 gain today, asked in March 2016; \$500 gain in three months, asked in March 2016; \$500 loss today, asked in March 2016; \$5000 gain today, asked in May 2016; \$2500 gain today, asked in May 2016; \$5000 gain today, asked in January 2017; \$5000 gain in three months, asked in January 2017; \$500 gain today, asked in January 2017; \$500 loss in three months, asked in January 2017; \$500 loss today, asked in March 2017; \$500 loss in three months, asked in March 2017; \$500 loss in two years, asked in March 2017.

and in which permanent earnings and wealth are continuous.

In the second specification (4.2), the term $perm_t^i$ is the continuous level of permanent earnings, the term $checking_t^i$ measures the amount that the head (and spouse if any) holds in checking and saving accounts, that is, the amount of fully liquid wealth. It is also a continuous variable. In this specification, the variables $perm_t^i$ and $checking_t^i$ are standardized: one unit corresponds to one standard deviation away from the mean.

I estimate the specifications described by (4.1) and (4.2) with linear regressions.

The effect of permanent earnings on the MPC. The first column in Table 1 presents the results from estimating specification (4.1). The first three lines in this column show that, holding non-housing wealth fixed, moving from the first to second and third quartiles of the permanent earnings distribution significantly raises the reported MPC of the respondents. Moving from the first to the second quartiles raises the MPC by 0.052, moving from the first to the third quartile raises the MPC by 0.045. There is no longer any significant effect of moving from the first to the fourth quartile. To put these values into perspective, note that the mean MPC in this sample is 0.176, so an increase by 0.05 corresponds to a 28% increase for someone at the mean MPC. The fact that raising permanent earnings raises the MPC is in line with the theoretical prediction that I obtain in the first section, but the effect is non-monotonic in this first specification: moving from the first to the second quartile of the permanent earnings distribution has a stronger effect than moving from the first to the third or fourth—although the differences are not statistically significant. However, this non-monotonicity disappears in subsequent specifications, in which I control for respondent characteristics. The 'Non-housing wealth Q' lines show that moving from the first to the third and fourth quartiles of the non-housing wealth distribution reduces the MPC, consistent with the existing theoretical and empirical literature.

The second column presents the results when I additionally control for respondent and household characteristics. The effect of moving from the first quartile of the permanent earnings distribution to all higher quartiles is positive and significant. The point estimates are now larger, in particular for moves to the third and fourth quartiles. The effect is also monotonic: the point estimates associated with higher quartiles are larger. In contrast, the effect of an increase in non-housing wealth is smaller.

The third column presents the estimation results among the subsample of singles, to whom the standard life-cycle model with a one-person household that I consider in the theoretical section most exactly applies. Indeed, in this subsample, the permanent earnings of the respondent coincides with the permanent earnings of the household, not with half of it, and there is only one decision maker in the household. The effect of moves to the top of the permanent earnings distribution on the MPC becomes larger: moving from the first to the third quartile of the distribution raises

	(1) MPC	(2) MPC	(3) MPC	(4) MPC	(5) MPC	(6) MPC	(7) MPC
Permanent earnings Q2	0.052** (0.021)	0.060*** (0.022)	0.031 (0.055)				
Permanent earnings Q3	0.045** (0.021)	0.063*** (0.023)	0.132** (0.063)				
Permanent earnings Q4	0.028 (0.022)	0.077*** (0.028)	0.188** (0.076)				
Permanent earnings (s.d.)				0.094** (0.043)			0.083** (0.041)
Permanent earnings (s.d.) ²				-0.010 (0.008)			-0.009 (0.007)
Non-housing wealth Q2	-0.011 (0.021)	0.002 (0.022)	0.018 (0.054)				
Non-housing wealth Q3	-0.046** (0.021)	-0.029 (0.022)	-0.040 (0.054)				
Non-housing wealth Q4	-0.064*** (0.022)	-0.036 (0.024)	-0.060 (0.064)				
Checking and saving accounts (s.d.)				-0.176** (0.075)	-0.117 (0.073)	-0.196*** (0.073)	
Checking and saving accounts (s.d.) ²				0.022 (0.017)	0.010 (0.017)	0.026 (0.017)	
Current earnings (s.d.)						0.075 (0.049)	
Current earnings (s.d.) ²						0.006 (0.021)	
Ratio (s.d.)							-0.828*** (0.287)
Ratio (s.d.) ²							0.146*** (0.050)
Constant	0.079*** (0.021)	-0.117 (0.071)	-0.152 (0.146)	0.026 (0.224)	0.007 (0.232)	-0.020 (0.225)	0.048 (0.228)
Question × date controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Respondent controls	No	Yes	Yes	Yes	Yes	Yes	Yes
Household controls	No	Yes	Yes	No	No	No	No
Singles only	No	No	Yes	Yes	Yes	Yes	Yes
Observations	2,733	2,733	493	250	252	252	250
R2	0.130	0.149	0.252	0.232	0.206	0.228	0.213

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1: The effect of permanent earnings on the MPC

the reported MPC by 0.132; moving from the first to the fourth quartile of the distribution raises the reported MPC by 0.188. This larger effect is consistent with the fact that single-household respondents are more exposed to their own earnings risk, since there is no second earner (or potential second earner) in their household.

The fourth column presents the results of estimating specification (4.2) among the subsample of single households, but with continuous variables and a more precise measure of liquid wealth. A one standard deviation increase in permanent earnings (\$46,217) around its sample mean raises the MPC by 0.094. The effect is statistically significant at the 5% level. It corresponds to a 50% increase at the mean MPC. The effect of the square of permanent earnings is smaller and not significant—recall that the effect of the level of permanent earnings on the MPC already captures a second cross-partial derivative of consumption, $\partial MPC_t / \partial e^{p_i} = \partial^2 c_t / (\partial a_t \partial e^{p_i})$, so the square captures a third cross partial derivative of consumption. Considering the effect of liquid wealth, one standard deviation increase in the amount held on checking and saving accounts (\$59,258) reduces the reported MPC by 0.176, significant at the 5% level. Comparing the two, the effect of permanent earnings is about half as large as that of liquid wealth. The effect of the square of liquid wealth is small and not significant. The fact that the effect of liquid wealth is larger in this specification than in the previous ones likely reflects the fact that the measure of liquid wealth more precisely captures truly liquid wealth. Both results are consistent with the theoretical predictions of the model.

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 appears modest absent controls for earnings. The first stylized fact that motivates the analysis is that some people with medium and high levels of liquid wealth still have a MPC that is significantly different from zero. This results in a modest estimated effect of liquid wealth on the MPC, since the MPC does not decrease very strongly as wealth reaches high levels. The positive effect of the permanent component of earnings on the MPC can account for this fact: in general people with higher levels of liquid wealth also have higher levels of permanent earnings. Although their higher liquid wealth reduces their MPC, their higher permanent earnings raises it. As a result, the MPC of people with high level of liquid wealth is smaller but not that much smaller than those with less liquid wealth. The related prediction is that the estimated effect of liquid wealth should be larger in absolute value when controlling for permanent earnings and eliminating its confounding effect than without such a control.

Column (5) presents quantitative evidence of this. It is based on an estimation of the same specification as (4) but in which I drop the permanent earnings regressors. In this new specification (5), the effect of liquid wealth is no longer significant. The point estimate of the effect is 0.12, about two thirds of its value in specification (4). Thus, failing to control for the level of permanent

earnings leads to finding a less significant and more modest effect of liquid wealth. The positive effect of permanent earnings can thus partly explain why existing studies find modest effects of liquid wealth.

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 an unexpected one-time income shock. Some studies even find a positive effect of current earnings on the MPC. This has been considered a surprising fact because one could expect the effect of current earnings to be similar to that of liquid wealth, since both provide additional immediately available resources. My result that permanent earnings raise the MPC can account for this second stylized fact. An increase in current earnings coming from its permanent component increases immediately available resources but increases risky future resources more, so it reduces the ratio of risk-free to risky resources. This raises the MPC. In contrast, an increase in current earnings coming from its more transitory component raises immediately available resources without raising risky resources and reduces the MPC. 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 is larger than that of liquid wealth and can be non-negative. It should be smaller than that of permanent earnings.

I confirm this stylized fact and quantify it in the SCE. Specifically, Column (6) in Table 1 shows that when I estimate a version of equation (4.2) in which I substitute the permanent component of annual earnings with total current annual earnings, its effect on the reported MPC is positive but not statistically significant. This is line with most empirical studies, finding a non-significant effect of current earnings on the MPC. The point estimate of the effect is smaller than that of permanent earnings.¹⁴

Does the ratio of liquid wealth over permanent earnings capture all of the effect of permanent earnings on the MPC? With isoelastic preferences and a zero consumption floor $c_0 = 0$, preferences are homothetic. In that case, permanent earnings do not affect the MPC beyond their effect through the ratio of liquid wealth over permanent earnings. In contrast, with a strictly positive consumption floor $c_0 > 0$ that generates a form of non-homotheticity, an increase in permanent earnings raises the MPC beyond its effect through the ratio. To test which of the two frameworks fits the data best, I examine the effect of permanent earnings on the MPC when controlling for

¹⁴Controlling for current earnings is however sufficient to capture the confounding effect of permanent earnings on liquid wealth: the impact of liquid wealth is significant and larger in this estimation (6) than in estimation (5), which corresponds to the case without any earnings controls.

the ratio of the amount held on checking and saving accounts over permanent earnings, instead of the amount on checking and saving accounts. The ratio is built with both variables expressed in \$ amount and is standardized.

Column (7) presents the results. The first order effect of a one standard deviation increase in this ratio around its mean reduces the reported MPC by 0.828. The effect is significant at the 1% level. The second order effect of the ratio is positive and significant. Combining both first and second order effects, the average effect of a one standard deviation increase at the mean is to decrease the MPC by 0.682. Note that, a one standard deviation of this ratio corresponds to a very large shock of seventeen times its mean value. Importantly, beyond this significant effect of the ratio of liquid wealth over permanent earnings, the level of permanent earnings still influences positively the MPC. This is consistent with the presence of a strictly non-zero consumption floor—or possibly with other non-homotheticities generating the same result. It also means that, when examining the determinants of the MPC, it is not sufficient to include the ratio of wealth over permanent earnings: one needs to include both wealth and permanent earnings—or both the ratio and permanent earnings.

Policy implication One policy implication is that, in order to maximize the consumption response to a fiscal stimulus, targeting the stimulus checks to people with low income is not the most effective choice. This is because people who have a low income because they have a low level of permanent earnings have a lower MPC everything else equal. I confirm this low effectiveness of targeting on income in the survey data if the objective is to maximize the response the stimulus. The average MPC in the bottom quartile of the earnings distribution is less than one point higher than the average MPC in the whole population: 0.183 versus 0.174.

A more effective strategy, based on the mechanism that I document, would be to use the permanent component of earnings combined with liquid wealth. However, permanent earnings is not typically observed by policy makers. As a substitute, I find that targeting based on a combination of age and liquid wealth is more effective than targeting based on income. The average MPC among people below median liquid wealth and below median age in my sample is more than five points higher than the average in the whole population: 0.225 versus 0.174. I detail these statistics in Table (??), Appendix ??.

4.2 Heterogeneity analysis, other tests of the theory, and robustness

Heterogeneity analysis. I break down the sample into MPCs out of positive versus negative shocks, small versus large shocks—equal to \$500 or larger than \$500—, and shocks realized today versus shocks announced today but realized in the future, typically in three months. Tables ??, ??,

and ?? in Appendix ?? present the results. The coefficients on the effect of permanent earnings on the MPC are larger and more often significant for positive than for negative shocks. This is true, in particular for the effect of moving to the top quartile of permanent earnings. The coefficients are also generally larger for smaller income shocks. The effects are markedly larger and more significant for shocks received today than in the future.

Robustness: controlling for overoptimism and overconfidence. Balleer, Duernecker, Forstner, and Goensch (2021) document an optimistic bias among SCE respondents with a low level of education and Pfäuti, Seyrich, and Zinman (2024) document that people with lower cognitive skills are overconfident about their future financial situation. The mechanism that I document implies that people with high expected future earnings also have a high MPC, because permanent earnings raises both expected future earnings and the MPC, but optimism and overconfidence may generate a similar correlation. To examine this point, I re-run the analyses adding a control dummy capturing whether people’s expectation at t of their earnings at $t + 1$ ended up higher than the realized earnings at $t + 1$. This dummy has a non-significant but positive effect on the MPC. Adding it makes little difference in the effect of permanent earnings on the MPC. Thus, optimism and overconfidence are not significant drivers of my findings. I present these results in Table ??, Appendix ??. The new dummy corresponds to the line ‘Expects higher earnings than realized’. These findings are not inconsistent with the presence of a meaningful amount of overoptimism and overconfidence. The optimism documented in Balleer, Duernecker, Forstner, and Goensch (2021) may be captured by the education dummies in the specification where I include them. In Pfäuti, Seyrich, and Zinman (2024), the first order mechanism is that, because overoptimistic people consume more than what would be their true optimum, they end up with lower wealth. They have a higher MPC mostly because of this lower wealth. Since I control for wealth, my mechanism would not capture this, but overconfidence may still substantially affect the dynamics of wealth accumulation.

Robustness: using a measure of consumption rather than of hypothetical MPC. 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 liquid wealth on consumption. Although the partial effect of liquid wealth is not exactly an MPC, since wealth is endogenous, I find that the interaction between the effect of liquid wealth and of permanent earnings on consumption is significant and positive. I present the methodology and results in Appendix ??.

5 Permanent earnings and the MPC in simulated data

To understand whether life-cycle models are not only qualitatively but also quantitatively consistent with these empirical findings, and whether they are able to reproduce them, I run numerical simulations.

5.1 Model and calibration

Consumers' maximization problem. I simulate and calibrate a standard incomplete market model that mimics the situation of US households. It follows closely the model of Kaplan and Violante (2010). 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. The period utility $u(\cdot)$ is a log-utility function. Consumers face a minimum yearly consumption threshold c_0 of \$ 1,638. The value is set to resemble the minimum yearly earnings threshold below which Guvenen, Karahan, Ozkan, and Song (2021) select people out.

Wealth. People only have access to one perfectly liquid asset in the model. The discount factor β is set to match a ratio of liquid wealth to annual earnings of 0.56, obtained from the Survey of Consumer Finance (SCF). This is the ratio that Kaplan and Violante (2022) use as a calibrating target when they compare how different versions of standard incomplete market models fare at matching MPC and wealth statistics. This means that I do a liquid wealth calibration: there is only one asset in the model, and it is liquid and calibrated to match the liquid wealth that people hold in the data. The underlying assumption is that people may hold illiquid wealth on top of their liquid wealth but they do not adjust their illiquid wealth to smooth consumption. Kaplan and Violante (2022) document that the average MPC obtained with such a liquid wealth calibration matches the data well, with a fit comparable to that of a two-asset model. This modeling choice does not make it possible to match or say anything about total wealth—while a two-assets model makes it possible—but total wealth is not the focus of this paper.

Discount factor. To obtain that the mean value of liquid wealth in the population equals 56% of the mean annual earnings in the population, I calibrate internally the baseline discount factor β . The mean earnings in my baseline model is \$59,964, so this implies a mean liquid wealth of \$33,601. The discount factor that matches it is $\beta = 0.9734$.

Interest rate. The yearly interest rate on the risk-free liquid asset is constant and set to $r = 0.01$, to match the low real interest rate on liquid holdings over the period 2016-2017.¹⁵

¹⁵The average 10-Year Real Interest Rate in the US over March 2016-March 2017 is 0.5% (see Federal Reserve

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 \$5,524 (in March 2016). This is coming from the SCF data about the credit card balance still owned (question x413). The top 90th percentile of unpaid balance is \$5,300 in the 2016 survey (conducted between March and December 2016). The top 90th percentile of unpaid balance is \$6,552 in the 2019 survey (conducted between March and December 2019). A rule of three implies a top 90th percentile over the survey period March 2016-March 2017 that I consider of \$5,524. That is the limit that I set.

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 91. I obtain the survival probabilities from Kaplan and Violante (2010), who use the life tables of the National Center for Health Statistics.¹⁶ If still alive at age 91, a household dies with certainty at age 92.

Earnings. My main departure from Kaplan and Violante (2010) is that I use a more general earnings process, which encompasses the simple transitory-permanent process that they use. I let 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 do 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). I set the parameters of this process equal to the estimates of Guvenen, Karahan, Ozkan, and Song (2021), summarized in Appendix ?? of this paper and taken from Table IV of their paper and Table D.III of their online appendix.

Taxes, transfers, and social security income. People pay taxes according to the nonlinear tax function of Gouveia and Strauss (1994), $tax(y_t^i) = \tau^b (y_t^i - ((y_t^i)^{-\tau^p} + \tau^s)^{-1/\tau^p})$ parametrized with $\tau^b = 0.258$, $\tau^p = 0.768$, $\tau^s = 2.0e - 4$ as in Kaplan and Violante (2010).¹⁷

Bank of Cleveland (2015-2018)). 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.

¹⁶See https://www.cdc.gov/nchs/products/life_tables.htm. Since, from the CDC, the period I consider is one with higher survival probabilities, I shift their probabilities, so that their survival from age 60 to age 61 is my survival from age 65 to age 66.

¹⁷Contrary 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

The government guarantees a minimum income of \$1,638 to people unemployed or with earnings below this threshold. The value is set to resemble the minimum yearly earnings threshold below which Guvenen, Karahan, Ozkan, and Song (2021) select people out. This means the government makes transfers of a value equal to the gap between earnings and this minimum threshold to individuals with labor earnings below this minimum threshold.

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. From this first bend point to a second bend point, it is 32 percent. It is 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 situations where individuals are hit by a one-time unexpected income shock. The first shock that I consider is a \$500 increase in beginning-of-period wealth. The second shock that I consider is a \$500 decrease in beginning-of-period wealth. I chose these two types of shocks because they are the two most common shocks in the survey: each make for 20% of the observations. I only simulate two different shocks instead of 13, for simplicity reasons. In the simulations, the unexpected income gain or loss occurs at random, once in the life-time, between age 26 and age 55. I simulate a gain and a loss for each consumer. They take place at the same age for each consumer. Similar to what I do in the survey data analysis, I pool together the MPCs out of both shocks, and I control for the effect of the type of MPC question in all analyses with a dummy capturing whether the shock is a \$500 gain or a \$500 loss.

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).¹⁸

Price harmonization. In the simulations, the income process is calibrated with the parameters estimated by Guvenen, Karahan, Ozkan, and Song (2021), but their estimation is based on data deflated and expressed in 2010\$ value. I simulate the model with their parameters, thus in 2010\$. I then convert the simulated values to express them in March 2016\$.

I use the tax function to recover net earnings.

¹⁸The number of grid points is as follows: the grid for wealth has 130 exponentially spaced grid points; the grid for the highly persistent component of earnings is age-varying and at each age has 21 equally spaced points; the grid for the transitory shock has 9 equally spaced points; the grid for the fixed effect component of earnings has 7 equally spaced points; the grid for lifetime average earnings (used to compute retirement income) has 11 equally spaced points. Expanding the grid further does not change the results in a noticeable way.

Building permanent earnings. In the simulations, I directly observe the fixed effect α^i and the highly persistent component of log-earnings p_t^i . Their product constitute the permanent component. I normalize it 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 log-earnings among employed people.

Selection. As in the empirical analysis, I select individuals below age 55 and employed at the moment when they experience the transitory shock.

	Survey data (SCE)	Simulated data
Average earnings	65,642	59,964
Non-housing wealth (converted from categories)	92,667	.
Checking and saving accounts*	26,274	.
Liquid wealth	.	33,601
Corr. permanent earnings categories/wealth categories	0.328	0.409
Share of people at the constraint	.	0.251
Yearly MPC	0.477	0.512
Observations	2,733	7,108

*The number of observations in the first column of this line is 1,143 (subset where the variable is observed).

Table 2: Model fit

Wealth and earnings comparison in the simulated and survey data. I compare simulated data to the survey data on statistics related to earnings, liquid wealth, and the MPC. The average earnings generated by the process and parameters in Guvenen, Karahan, Ozkan, and Song (2021) is close to, although a little lower than, 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 (members contributing to the rent or owning the house), who might earn a slightly higher wage than non-heads, while the administrative data used in Guvenen, Karahan, Ozkan, and Song (2021) covers the annual earnings of all.

The level of liquid wealth in the simulations is targeted to 56% of the average annual earnings, which corresponds to \$33,601. This target is based on the ratio observed in the SCF. It is not based on the survey data that I use. Yet the liquid wealth reported by survey respondents is not too different from the liquid wealth obtained by multiplying the SCF ratio with the average annual earnings implied by Guvenen, Karahan, Ozkan, and Song (2021)'s earnings process. The average amount on checking and saving accounts in the subset of respondents for whom this value is observed is \$26,274. Also, when I convert the non-housing wealth categories of the SCE into

wealth levels by attributing to each respondents the medium level of wealth in their categories—also \$250 for those who answer less than \$500 and 1,5 million for those who answer more than 1 million—, I find that the average non-housing wealth is \$92,667. The definition of non-housing wealth includes categories of wealth not considered to obtain the SCF ratio. Thus, the average liquid wealth in the numerical simulations lies in between the most and least restrictive measures of liquid wealth in the SCE, close the most restrictive one, that is, the amount on checking and saving accounts.

The correlation between the permanent earnings quartiles and the fourteen categories of non-housing wealth—which I construct in the simulations as well—is positive but not too large in both the survey and the simulated data. It is lower in the survey data, at 0.328, than in the simulations, at 0.409.

In the simulated data, a little more than one quarter of the people have their liquid wealth at the minimum possible level of -\$5,524. This share is not far from 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 0.251, with a range going from 0.220 to 0.503.

Finally, the yearly MPC in the model is close to the yearly MPC implied by the survey data.

5.3 Specification and main results

Specification. With these simulated data, I estimate a specification close to (4.2). The equation that I estimate is

$$MPC_t^i = c_1 + c_2 perm_t^i + c_3 (perm_t^i)^2 + c_4 wealth_t^i + c_5 (wealth_t^i)^2 + c_6 X_t^i + \xi_t^i. \quad (5.1)$$

The controls X_t^i are dummies for the type of MPC question—positive or negative shock—and age categories—the only demographic characteristic in the simulated data is age. The age categories are the same as in the empirical estimation.

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 three columns are a reminder of the results I obtain in the survey data when estimating (4.2), the survey counterpart to (5.1).

Column (4) shows that, in the numerical simulations, an increase in permanent earnings raises the MPC when controlling for wealth. The coefficient implies that a one standard deviation increase in the permanent component of earnings raises the MPC by 0.182. This is close to the increase by 0.220 estimated in survey data. The effect of liquid wealth is negative. The coefficient implies that a one standard deviation increase in liquid wealth reduces the MPC by 0.327. This is

	Survey data			Simulated data (baseline)		
	(4)	(5)	(6)	(4)	(5)	(6)
Yearly MPC						
Permanent earnings (s.d.)	0.220** (0.100)			0.182		
Permanent earnings (s.d.) ²	-0.024 (0.018)			-0.007		
Checking and saving account (s.d.)	-0.410** (0.175)	-0.273 (0.169)	-0.457*** (0.169)	-0.327	-0.200	-0.300
Checking and saving account (s.d.) ²	0.052 (0.040)	0.023 (0.040)	0.062 (0.039)	0.008	0.005	0.007
Current earnings (s.d.)			0.174 (0.114)			0.146
Current earnings (s.d.) ²			0.015 (0.049)			-0.006
Constant	0.060 (0.522)	0.017 (0.541)	-0.048 (0.524)	0.501	0.509	0.539
Average yearly MPC	0.591	0.592	0.592	0.512	0.512	0.512
Question $\times$ date controls	Yes	Yes	Yes	Yes	Yes	Yes
Respondent controls	Yes	Yes	Yes	Yes	Yes	Yes
Household controls	No	No	No	No	No	No
Singles only	Yes	Yes	Yes	Yes	Yes	Yes
Observations	250	252	252	7,108	7,108	7,108

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3: The effect of permanent earnings on the MPC in survey data and in simulations

similar to, although a little smaller than, the decrease by 0.410 estimated in survey data.

Columns (5) and (6) show that the simulations are able to reproduce the two stylized facts that motivate the analysis. First, when failing to control for permanent earnings, the estimated effect of liquid wealth drops by about one third of its value: the coefficient associated with wealth move from -0.327 to -0.200. Second, the effect of current earnings on the MPC is positive. It thus differs from the effect of liquid wealth, which is negative. The coefficient on the effect of current earnings is 0.146. It is lower than that of permanent earnings, in line with the fact that current earnings is made of two components, a permanent component, which raises the MPC, and a transitory component, which reduces it.

5.4 The role of the rich earnings process: why can the model generate the stylized facts?

My simulations are able to generate coefficients comparable to those I estimate in survey data and to account for the two stylized facts that motivate the analysis. The model that I use is standard, except for my incorporating the rich earnings process proposed in Guvenen, Karahan, Ozkan, and Song (2021). I thus examine what elements of the model are important for matching the data, and in particular to what extent this richer earnings process, or some of its elements, matters.

Examining the role of the model components To examine the role of the different elements of the model, I start from a simple framework and add the elements of the baseline model one by one. The simple framework includes a simple transitory-permanent earnings specification. There is no non-employment shocks, the transitory and highly persistent shocks are drawn from normal distributions rather than mixture of normals,¹⁹ and there is no variance in individual fixed effects α —although the initial value of the highly persistent component is drawn with non-zero variance. In this simple framework, I also make social security income proportional to the level of permanent earnings on the last year of work, instead of having a progressive social security income. The proportionality coefficient is such that the average pension income is the same as in the baseline. I also remove the exogenous borrowing constraint. This corresponds to Model 1. This model encompasses the simple one that I present in the theoretical section of the paper—except for having $\rho = 0.991$ in the simulations and $\rho = 1$ in the theoretical section. In Model 2, I reintroduce the progressive social security income. In Model 3, I additionally reintroduce the exogenous borrowing constraint. In Model 4, I additionally reintroduce that transitory and highly persistent shocks are drawn from a mixture of normal distributions. In Model 5, I additionally reintroduce non-employment shocks. Finally, additionally reintroducing a variance in the fixed effect α yields back the baseline model.

In each model, the selection is the same as in the baseline. Thus, when there are no non-employment shocks and no highly persistent shocks drawn from extreme distributions, the sample is larger. Indeed, there are no non-employed individuals and no individuals whose earnings are low enough to receive the transfer, which are those I select out.

The effect of permanent earnings with and without the rich earnings process Each column in each panel of Table 4 presents the results of specification (4), specification (5), and specification (6). These are the same specifications I consider in Table (3), except I present them by rows rather

¹⁹The variance of the normal distributions are those of the most likely normal distribution in the mixture specification. The fixed effect is adjusted so the average level of earnings is the same as in the Guvenen, Karahan, Ozkan, and Song (2021) specification.

	Model 1 Most simple	Model 2 1+ progressive SSI	Model 3 2+ borrowing constraint
Calibrated β	0.9781	0.9817	0.9709
<i>Specification (4)</i>			
Permanent earnings (s.d.)	0.0002	-0.0002	-0.0573
Permanent earnings (s.d.) ²	-0.0002	-0.0000	0.0064
Checking and saving account (s.d.)	-0.0083	0.0009	0.0481
Checking and saving account (s.d.) ²	0.0018	-0.0001	-0.0029
<i>Specification (5)</i>			
Checking and saving account (s.d.)	-0.0081	0.0007	-0.0203
Checking and saving account (s.d.) ²	0.0018	-0.0001	0.0031
<i>Specification (6)</i>			
Current earnings	-0.0000	-0.0002	-0.0632
Current earnings ²	0.0001	-0.0000	0.0004
Average MPC	0.041	0.039	0.455
Observations	10,000	10,000	10,000
	Model 4 3+ mixture	Model 5 4+ non-employment	Baseline 5+ fixed effect
Calibrated β	0.9714	0.9732	0.9734
<i>Specification (4)</i>			
Permanent earnings (s.d.)	0.2883	0.2393	0.1820
Permanent earnings (s.d.) ²	-0.0355	-0.0177	-0.0074
Checking and saving account (s.d.)	-0.2350	-0.3504	-0.3267
Checking and saving account (s.d.) ²	0.0149	0.0123	0.0084
<i>Specification (5)</i>			
Checking and saving account (s.d.)	-0.0208	-0.1782	-0.1996
Checking and saving account (s.d.) ²	0.0004	0.0040	0.0050
<i>Specification (6)</i>			
Current earnings	0.2634	0.1961	0.1462
Current earnings ²	-0.0304	-0.0134	-0.0056
Average MPC	0.387	0.506	0.512
Observations	9,980	7,108	7,108

Table 4: The effect of permanent earnings on the MPC in the simulations of different models

than by columns and show a more limited set of coefficients. The first column corresponds the most simple model. The results of specification (4) show that permanent earnings have a positive though very small first order effect on the MPC. This is consistent with the prediction of the theoretical section: an increase in permanent earnings raises the MPC. However, with a simple earnings process, the amount of earnings risk is limited, and the effect of permanent earnings on the MPC remains very small. The effect of liquid wealth is negative. It is smaller than in the data but less so than the effect of permanent earnings. The results in specification (5) show that, when the effect of permanent earnings on the MPC is very small, controlling for permanent earnings barely changes the estimated effect of wealth on the MPC: it moves from -0.0083 with a control to -0.0081 without a control. Finally, the results in specification (6) show that, when the effect of permanent earnings on the MPC is very small, the effect of an increase in current earnings on the MPC is closer to that of an increase in transitory earnings, which is negative. As a result, although permanent earnings has a positive effect on the MPC, this model cannot reproduce the two stylized facts that motivate this paper. The average MPC in this model is very small.

The second column of the top panel presents the results when introducing the progressive social security income scheme to Model 1. The coefficients all remain very small. Several of them change sign. This suggests that making social security income progressive is not what drives my result of a large and positive effect of permanent earnings on the MPC. If anything, it has the opposite effect of making people with higher permanent earnings slightly less sensitive to one-time income shock. This is consistent with the fact that a more progressive system provides insurance and reduces the need for precautionary saving. The MPC remains very small.

The third column of the top panel presents the results when adding a \$5,524 exogenous borrowing limit to Model 2. This model is close to the one simulated for instance in Kaplan and Violante (2010) or to the one with a liquid wealth calibration in Kaplan and Violante (2022). From specification (4), permanent earnings have a small negative effect on the MPC and liquid wealth a small positive effect on the MPC. This indicates that the exogenous borrowing constraint alone is not what drives my results of a large and positive effect of permanent earnings on the MPC: introducing an exogenous budget constraint without a richer earnings process pushes the results further away from the empirical findings. When estimated alone, in specification (5), the effect of liquid wealth becomes negative, as observed empirically. It remains modest. From specification (6), the effect of current earnings on the MPC is small and negative. Thus, this model cannot reproduce the two stylized facts that motivate this paper. The average MPC produced by this model is fairly large, at 0.455. These results are consistent with the overall findings of this class of models: an increase in liquid wealth, typically measured without permanent earnings controls, reduces the MPC but the average effect is relatively modest; this is true although the MPC can be quite large.

The first column of the bottom panel presents the results when adding to Model 3 that the tran-

sitory and highly persistent earnings shocks are drawn from the mixture of normal distributions estimated in Guvenen, Karahan, Ozkan, and Song (2021) rather than simply from normal distributions. Note that the mixtures are such that consumers have a relatively high probability to draw their shocks from a distribution with the typical variance but there is a small probability to draw from a distribution with a high variance. The results from specification (4) show that adding this element turns the effect of permanent earnings positive and that of liquid wealth negative, and make both much larger in absolute value: a one standard deviation increase in permanent earnings raises the yearly MPC out of a \$500 gain by 0.288; a one standard deviation reduces the yearly MPC out of a \$500 gain by 0.235. Both coefficients are much closer to the empirical estimates, although the effect of permanent earnings is now a little too high and that of liquid wealth a little too low. This model is able to generate qualitatively the two stylized facts that motivate the analysis: the effect of liquid wealth appears much more modest when not controlling for the level of permanent earnings (specification (5)), and an increase in current earnings raises the MPC (specification (6)). This suggests that, because people face substantially larger earnings risk, an increase in permanent earnings that amplifies this risk exposure modifies their behavior much more strongly. The average MPC drops a little compared to Model 3. This might be reflecting the interaction between liquidity constraint and precautionary motive: since the precautionary motive is stronger, people hold more liquid wealth and the borrowing constraint is less binding on average.

The second column of the bottom panel presents the results when adding non-employment shocks to Model 4, according to the distribution and parameters estimated in Guvenen, Karahan, Ozkan, and Song (2021). Specification (4) presents coefficients that are very close to the one that I estimate empirically in survey data. Specifications (5) and (6) also yield results that are quantitatively similar to the estimation of similar specifications in survey data. This suggests that the presence of non-employment shocks are important to match more closely the negative effect of liquid wealth that I estimate in survey data. It also yields a higher average MPC, closer to the one that respondents report in survey data.

Finally, the third column of the bottom panel reproduces the baseline results already shown in the left part of Table 3. They are close to the results in Model 5. This suggests that the variance of fixed effects does not play a very important role in the simulations.

6 Conclusion

In this paper, I establish analytically that, in the standard life-cycle model, everything else equal, consumers with a higher permanent component of earnings have a higher MPC. The mechanism is through precautionary behavior: everything else equal, consumers with a higher permanent component of earnings have more resources but a larger share of their resources are risky. These people

consume more, because they have more resources, but save a larger fraction of their expected lifetime resources for precautionary reasons, because their resources are on average more risky. Their consumption is more constrained by the uncertainty about future earnings: it would increase more than that of others if the uncertainty disappeared. A windfall gain that increases the share of certain, realized resources in total expected resources leads them to reduce this precautionary saving more and thus consume more out of the windfall gain.

I find that this theoretical prediction holds true in the New York Fed Survey of Consumer Expectations. First, among employed workers who share the same demographic characteristics, expected future earnings can constitute a proxy for permanent earnings. I verify that an increase in this proxy for permanent earnings raises linearly the standard deviation of future earnings. This means that, in the data, an increase in the permanent component of earnings is an increase in risky resources in the same way as under a transitory-permanent earnings specification.

Second, controlling for wealth and demographics, this proxy for the permanent component earnings covaries positively with the respondents' MPCs. A one-standard deviation increase in the permanent component of earnings associates with a 0.09 increase in the reported quarterly MPC out of a hypothetical one-time income shock. In absolute value, the effect of permanent earnings is of a comparable magnitude to that of liquid wealth. More precisely, it is between one-third and two-third as large. Since permanent earnings and liquid wealth are positively correlated, while their effects on the MPC go in opposite directions, the effect of permanent earnings can offset a substantial part of the effect of liquid wealth on the MPC when it is not controlled for. This can explain why empirical studies typically find only a modest effect of liquid wealth on the MPC—or relatedly why the MPC does not fall that much with wealth and can remain far from zero even at high levels of liquid wealth.

I then show that these empirical findings are 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 MPC is positive and about as large as the one that I estimate in survey data. The effect of liquid wealth is negative and close to, although a little smaller than, the one that I estimate in survey data. The MPC level is close to the one that I observe in survey data. The numerical simulations can reproduce the empirical observations that failing to control for permanent earnings leads to underestimating the effect of liquid wealth, and that current earnings do not have a negative effect on the MPC. Incorporating the realistic and rich earnings process of Guvenen, Karahan, Ozkan, and Song (2021) is key to match quantitatively the empirical observations, in particular the fact that earnings innovations are drawn from a mixture of normal distributions and the fact that people are subject to non-employment shocks. The reason is that these features generate more earnings risk, which bolsters the precautionary motive and thus the magnitude of the mechanism that I identify.

References

- Aguiar, Mark A, Bils, Mark, and Boar, Corina (2020).** *Who Are the Hand-to-Mouth?* Working Paper 26643. National Bureau of Economic Research.
- Arellano, M., Attanasio, O., Crossman, S., and Sancibrián, V. (2024).** *Estimating Flexible Income Processes from Subjective Expectations Data: Evidence from India and Colombia.* Working Paper 32922. National Bureau of Economic Research.
- Arellano, M., Attanasio, O., De Nardi, M., Borella, M., and Paz-Pardo, G. (2024).** *Subjective Earnings and Employment Dynamics.* Tech. rep. mimeo.
- Arellano, M., Bonhomme, S., Vera, M. De, Hospido, L., and Wei, S. (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., 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.*
- 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.
- Baugh, Brian, Ben-David, Itzhak, Park, Hoonsuk, and Parker, Jonathan A. (2021).** “Asymmetric Consumption Smoothing.” *American Economic Review* 111.1, 192–230.
- Bilbiie, Florin, Galaasen, Sigurd M, Gürkaynak, Refet S., Mæhlum, Mathis, and Molnar, Krisztina (2025).** *HANKSSON.* Tech. rep. mimeo.
- 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 (2025).** “Five Facts about MPCs: Evidence from a Randomized Experiment.” *American Economic Review* 115.1, 1–42.
- Borusyak, K., Jaravel, X., and Spiess, J. (2024).** “Revisiting Event-Study Designs: Robust and Efficient Estimation.” *The Review of Economic Studies* 91.6, pp. 3253–3285. ISSN: 0034-6527.
- 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.
- Bunn, Philip, Le Roux, Jeanne, Reinold, Kate, and Surico, Paolo (2018).** “The consumption response to positive and negative income shocks.” *Journal of Monetary Economics* 96, pp. 1–15. ISSN: 0304-3932.
- 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.
- Christelis, D., Georgarakos, D., Jappelli, T., Pistaferri, L., and Rooij, M. van (2019).** “Asymmetric Consumption Effects of Transitory Income Shocks.” *Economic Journal* 129.622, pp. 2322–2341.
- Colarieti, Roberto, Mei, Pierfrancesco, and Stantcheva, Stefanie (2024).** *The How and Why of Household Reactions to Income Shocks*. Working Paper 32191. National Bureau of Economic Research.
- Commault, Jeanne (2025).** “Precautionary Saving At the Top.” 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.
- Federal Reserve Bank of Cleveland (2015-2018).** *10-Year Real Interest Rate [REAINTRATREARAT10Y]*.
- Federal Reserve Bank of New York (2015-2019).** *Survey of Consumer Expectations*.
- Friedman, Milton (1957).** *A Theory of the Consumption Function*. Princeton University Press.
- 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.
- Golosov, Mikhail, Graber, Michael, Mogstad, Magne, and Novgorodsky, David (Oct. 2023).** “How Americans Respond to Idiosyncratic and Exogenous Changes in Household Wealth and Unearned Income.” *The Quarterly Journal of Economics* 139.2, pp. 1321–1395. ISSN: 0033-5533.
- Gouveia, Miguel and Strauss, Robert P. (1994).** “Effective Federal Individual Tax Functions: An Exploratory Empirical Analysis.” *National Tax Journal* 47.2, pp. 317–39.
- Graziani, G., Klaauw, W. van der, and Zafar, B. (2016).** “Workers’ Spending Response to the 2011 Payroll Tax Cuts.” *American Economic Journal: Economic Policy* 8.4, 124–59.
- Güvenen, 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.
- Güvenen, Fatih, Ozkan, Serdar, and Madera, Rocio (2024).** “Consumption dynamics and welfare under non-Gaussian earnings risk.” *Journal of Economic Dynamics and Control* 169. Heterogeneity and Macroeconomics of Labor Markets, p. 104945.
- Hall, R. (1978).** “Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence.” *Journal of Political Economy* 86.6, pp. 971–987.
- Jappelli, Tullio and Pistaferri, Luigi (2014).** “Fiscal Policy and MPC Heterogeneity.” *American Economic Journal: Macroeconomics* 6.4, pp. 107–36.
- Johnson, D., Parker, J., and Souleles, N. (2006).** “Household Expenditure and the Income Tax Rebates of 2001.” *American Economic Review* 96.5, pp. 1589–1610.
- 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.

- Kotsogiannis, C. and Sakellaris, P. (2025).** *MPCs Estimated from Tax Lottery, Survey and Administrative Data*. Tech. rep.
- Koşar, G. and Melcangi, D. (2025).** *Subjective Uncertainty and the Marginal Propensity to Consume*. Tech. rep.
- 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.
- Maxted, Peter, Laibson, David, and Moll, Benjamin (Sept. 2024).** “Present Bias Amplifies the Household Balance-Sheet Channels of Macroeconomic Policy*.” *The Quarterly Journal of Economics* 140.1, pp. 691–743. ISSN: 0033-5533.
- 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, J., Ramey, V., and Wieland, J. (Feb. 2025).** *Micro MPCs and Macro Counterfactuals: The Case of the 2008 Rebates*. Tech. rep., qjaf015.
- Parker, J. (1999).** “The Reaction of Household Consumption to Predictable Changes in Social Security Taxes.” *The American Economic Review* 89.4, pp. 959–973.
- Parker, J. and Souleles, N. (2019).** “Reported Effects versus Revealed-Preference Estimates: Evidence from the Propensity to Spend Tax Rebates.” *American Economic Review: Insights* 1.3, pp. 273–90.
- 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.
- Pfäuti, Oliver, Seyrich, Fabian, and Zinman, Jonathan (2024).** “Bad luck or bad decisions? Macroeconomic implications of persistent heterogeneity in cognitive skills and overconfidence.” *mimeo*.
- Pistaferri, Luigi (2001).** “Superior Information, Income Shocks, And The Permanent Income Hypothesis.” *The Review of Economics and Statistics* 83.3, pp. 465–476.
- Rozsypal, F. and Schlafmann, K. (2023).** “Overpersistence Bias in Individual Income Expectations and Its Aggregate Implications.” *American Economic Journal: Macroeconomics* 15.4, pp. 331–71.
- Shapiro, M. and Slemrod, J. (2009).** “Did the 2008 Tax Rebates Stimulate Spending?” *American Economic Review* 99.2, 374–79.
- Straub, L. (2019).** “Consumption, Savings, and the Distribution of Permanent Income.” *mimeo*.
- 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.