Acknowledgement: the author declares that she has no relevant or material financial interests that relate to the research described in this paper.

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

Jeanne Commault*

March 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. This can explain two documented stylized facts: (i) the extent to which MPCs decrease with liquid wealth is modest enough that MPCs remain high among people with high levels of liquid wealth; (ii) MPCs do not decrease with current earnings although current earnings increase immediate liquidity. In a standard consumption model, permanent earnings raise the MPC because they raise risky future earnings without raising risk-free liquid wealth, strengthening precautionary behavior. In survey data, MPCs increase with permanent earnings when controlling for wealth and demographics. The effect is large enough to explain the two stylized facts. Numerical simulations can replicate survey results and generate the two stylized facts.

Key words: Marginal Propensity to Consume, 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. MPC distributions matter for several economic questions including the transmission of economic fluctuations, the design of stabilization

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

policies, and the dynamics of wealth accumulation. The new insights brought by heterogeneous agent models rely in important ways on the MPC heterogeneity in these models (see Kaplan, Moll, and Violante 2018).

An empirical result that emerged over the past decade is that, on average, people respond significantly to a one-time unexpected shock. 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 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, while only few people hold no wealth at all, a large enough fraction of people hold no liquid wealth to generate a non-zero average MPC. Other papers introduce different preferences and behavioral biases, such as temptation and present bias and generate a non-zero average MPC. The mechanisms behind also typically involve having more people with little liquid wealth and a high MPC, generally because some consumers have a lower desire to save.

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 remains modest, so some people with medium or high levels of liquid wealth have non-zero MPCs; and (ii) higher current earnings do not associate with lower MPCs, although people with higher current earnings have more resources and could be less constrained in terms of liquidity; higher current earnings can even associate with higher MPCs everything else being equal. These two facts also hold for MPCs associated with nondurable spending.

In this paper, I show that a standard consumption model can account for those two stylized facts when earnings have a permanent component—or a fixed effect component—that multiplies the realizations of future earnings. Indeed, in this model, an increase in the permanent component of earnings raises the MPC because it raises exposure to future earnings risk. At the same level of liquid wealth, consumers with higher permanent earnings have a stronger precautionary motive, and are thus more responsive to a windfall of liquidity. This positive effect of permanent earnings can explain fact (i): since liquid wealth and permanent earnings correlate, without a control for permanent earnings, the effect of liquid wealth is underestimated and appears more modest than it is; consumers with both high wealth and high permanent earnings do not have a much lower MPC than consumers with lower wealth and lower permanent earnings. The positive effect of permanent earnings can also explain fact (ii): current earnings, which includes a permanent component, may not

negatively affect the MPC when the positive effect of permanent earnings dominates that of the other components of current earnings. I provide survey evidence consistent with this mechanism. I show that numerical simulations of the standard model can quantitatively account for the survey evidence and generate the two stylized facts.

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, 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 uncertain thus risky. At the same level of liquid wealth, people with a higher level of permanent earnings have then a higher share of risky resources in their total expected resources. They consume more but save a larger fraction of their total expected resources for precautionary reasons. Because a one-time income gain reduces the need for saving by providing immediate liquidity, it has a stronger dissaving effect on consumers who have a higher wealth target before the income gain. Thus, everything else equal, consumers with higher permanent earnings reduce their saving more and raise their consumption more in response to a one-time income gain. For the same reason, they are also more responsive to a one-time income loss.

The result holds with homothetic preferences. In that case, the MPC is determined by the ratio of wealth to permanent earnings. A consumer with permanent earnings twice as large and a wealth that is either the same or larger but less than twice as large than another consumer has a lower wealth-to-permanent earnings ratio and thus a higher MPC. Note that the model does not imply that wealth and permanent earnings should always be proportional. At the same initial wealth, because people with higher permanent earnings save more, their wealth grows larger over time but their wealth-to-permanent earnings ratio remains lower. It is only asymptotically that in an infinite horizon model without shocks the wealth-to-permanent earnings ratio of the two consumers will be the same. The result still holds when consumers additionally face a Stone-Geary consumption floor. In that case, an increase in permanent earnings raises the MPC even when holding the ratio of wealth to permanent earnings constant: a consumer with both permanent earnings and wealth twice as large than another consumer has a higher MPC.

This positive effect of permanent earnings on the MPC can explain the stylized facts (i) and (ii) discussed above. 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 confounds 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 other less persistent components, may not reduce the MPC if the positive effect of the permanent component offsets the negative effect of the less persistent 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 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). In line with the presence of a permanent component that multiplies the future possible realizations of earnings, the variance of future earnings increases with the expected value of future 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 fraction of the effect of liquid wealth, and that people with medium to high levels of liquid wealth still have non-zero MPCs, 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 should also be able to offset the effect of the transitory component of earnings and obtain fact (ii). Consistent with this, the coefficient associated with current earnings is positive but smaller than that of permanent earnings. One policy implication of this result is that, in order to maximize the consumption response, targeting stimulus checks to people based on their income is not the most effective. This is because income includes components that both raise and reduce the MPC. Targeting based on a mix of wealth and age is more

effective. I find evidence that permanent earnings affect the MPC beyond the effect of the permanent earnings to wealth ratio. This is consistent with the existence of a non-zero consumption floor à la Stone-Geary.

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. More precisely, when I shift away from their process to the commonly used transitory-permanent process with normally distributed shocks, the effect of permanent earnings on the MPC is much smaller. Simulations with this simple process are also unable to generate the stylized facts (i)-(ii). This might be why the positive effect of permanent earnings on the MPC has been overlooked.

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 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 became an important 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: even households with medium to high levels of liquidity have a non-zero response to a one-time shock. 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).[1] The stylized fact (ii) is that, in contrast to that of liquid wealth, the effect of total current earnings on the

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

MPC is typically not significant. This is one of the results in 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), 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)[2] and of Lewis, Melcangi, and Pilossoph (2022). In Jappelli and Pistaferri (2014), the authors document a significant and negative correlation between current income and the MPC out of a shock proportional to people's income, which is 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) present new stylized facts on the relation between the MPC and the standard deviation of earnings growth, while I consider 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 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, although not directly its effect on the MPC. Among them, Carroll (2006) notes that one can normalize the standard consumer's problem by the permanent component of earnings to solve it with one less variable. This is useful in particular when simulating such a model. I expand the insight to show that the consumption solution is even homogeneous of degree one in wealth and permanent earnings,[3] and that the homogeneity has implications for the impact of permanent earnings on the MPC. 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) show the homogeneity of consumption 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 the OLG version of a consumption model with discrete time, where agents are born with zero wealth and the permanent component of earnings

[2]The 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.

[3]This is similar but 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.

is a time-invariant fixed effect (his Proposition 1). His paper then examines the effect of permanent earnings on this MPCP, that is, the concavity of consumption in permanent earnings. I prove homogeneity in a similar discrete-time setting, but letting initial wealth differ from zero and permanent earnings be subject to shocks. I also consider a departure from homothetic preference by introducing a consumption floor (à la Stone-Geary). In that case, the MPC increases when wealth and and permanent earnings are multiplied by the same strictly positive value (while under homogeneity it would stay the same). This latter result still implies that consumption increases strictly with the permanent component of 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, including an increase the average MPC generated by the model. I also find that this increases the MPC. Additionally, I show that introducing this process makes permanent earnings affect MPC more strongly. This 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 in a standard framework, I consider a simple version of the income fluctuations problem. There is a transitory-permanent earnings process and only one asset.

A consumer i is finite-lived, with T the length of their life. The consumer chooses consumption expenditures at period t, denoted c_t^i, to maximize lifetime expected utility subject to restrictions on the utility and earnings process, a budget constraint, and a terminal

condition on wealth

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

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

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

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

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

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

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 $u(.)$ is isoelastic with $\rho \geq 0$ such that $u'(c) = c^{-\rho}$. When $\rho = 1$, then $u(.) = ln(.)$. Isoelastic utility functions constitute the only class of utility functions with constant relative risk aversion (CRRA) so one can equivalently think of the utility as 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 $e^{p_t^i}$ that evolves as a multiplicative random walk and of a transitory innovation $e^{\varepsilon_t^i}$ that is an i.i.d. shock. Because the permanent component at t $e^{p_t^i}$ multiplies the value of the permanent component at the next period $t+1$, which itself multiplies the permanent component at $t+2$, it multiplies each future realization of earnings until the rest of the consumer's lifetime: at $t+s$, earnings are

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

The permanent component of earnings $e^{p_t^i}$ 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 $e^{p_t^i} = e^{p^i}$. This is for instance the specification in Straub (2019). Throughout this paper, it is not necessary that consumers face permanent income shocks e^{η}, but simply that they face some uncertainty about their future income so they have a precautionary saving motive. Uncertainty coming only from future transitory shocks is sufficient. I thus only impose that people face a strictly positive amount of transitory earnings uncertainty: $var(\varepsilon) > 0$ and $var(\eta) \geq 0$. 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.

The budget constraint (2.5) states that, to store their wealth from one period to the next the consumer only has access to one risk-free liquid asset. The term a_t^i denotes the level of this asset at the beginning of period t—or at the end of $t-1$. The risk-free return rate is r.

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. Thus, 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

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

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

MPC definition. What I want to examine is the effect of permanent earnings on the MPC, where the MPC is defined as the response of consumption to an unexpected one-time income shock such as a fiscal stimulus check. 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 the beginning-of-period wealth term 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 a small 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 is easy to obtain because it builds on the combination of two existing results. First, with isoelastic utility and no consumption floor, Carroll (2006) notes 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 in the case without a consumption floor, I further prove that the consumption solution of the consumer's problem is not just scalable in permanent earnings but homogeneous of degree one in both wealth and permanent earnings (Lemma 4, Appendix **??**). I do this because it helps understand the case with a consumption floor.[4]. The consumption policy function is $c_t = f^t(a_t, e^{p_t}, e^{\varepsilon_t})$. The homogeneity of degree one in a_t and e^{p_t} implies $c_t/e^{p_t} = f^t(a_t/e^{p_t}, 1, e^{\varepsilon_t})$: if one divides both a_t and e^{p_t} by e^{p_t}, the optimal level of consumption is also divided by e^{p_t}. The homogeneity of degree one of c_t implies an homogeneity of degree zero of the partial derivatives of c_t, including the MPC.

[4]Scalability is not exactly the same as homogeneity of degree one. Under scalability, if $c_t = f^t(a_t, e^{p_t}, e^{\varepsilon_t})$, then there exist a policy function $\tilde{f}^t$ with the same properties as f^t such that $c_t/e^{p_t} = \tilde{f}^t(a_t/e^{p_t}, 1, e^{\varepsilon_t})$. Yet the parameters of the new policy function $\tilde{f}^t$ may not be the same as that of the original policy function. Under homogeneity of degree one in both a_t and e^{p_t}, the two policy functions coincides with $c_t/e^{p_t} = f^t(a_t/e^{p_t}, 1, e^{\varepsilon_t})$. Although the fact that consumption is homogeneous of degree one in a_t and e^{p_t} did not sound new, I could not find it proved in my setting.

Indeed, when I differentiate $c_t = e^{p_t} f^t(a_t/e^{p_t}, 1, e^{\varepsilon_t})$ with respect to a_t I have

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

where $f_1(.,.,.)$ is the partial derivative of the policy function with respect to its first argument. There is homogeneity of degree zero: 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 with utility function displaying Hyperbolic Absolute Risk Aversion (HARA)—a set that includes all functions displaying constant relative risk aversion— and no consumption floor, consumption is strictly concave in liquid wealth. For any a_t, e^{p_t}, and e^{ε_t}

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

where $f_{11}(.,.,.)$ is the second-order partial derivative of the policy function with respect to its first argument. I further show in Commault (2024) that, for the subcase of HARA functions that are isoelastic, the concavity comes from precautionary saving decreasing more with wealth at lower levels of wealth. Non-isoelastic HARA utility functions generate concavity through other mechanisms. Combining these two results implies that the MPC is increasing in e^{p_t} when $a_t > 0$. To show this, I differentiate both sides of (2.8) with respect to e^{p_t}. Because the MPC only depends on the ratio e^{p_t}/a_t, the effect of e^{p_t} on the MPC writes as the opposite of the effect of a_t, multiplied by a_t/e^{p_t}. Because the MPC decreases with a_t, then it must increase with e^{p_t}

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

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) and a decrease in wealth that puts it back to its initial level (which raises the MPC). Thus, an increase in permanent earnings raises the MPC.

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 then a decrease in debt (which reduces the MPC). In that case, having a higher level of permanent earnings has the same impact on the MPC as having more wealth, not less wealth, so it reduces the

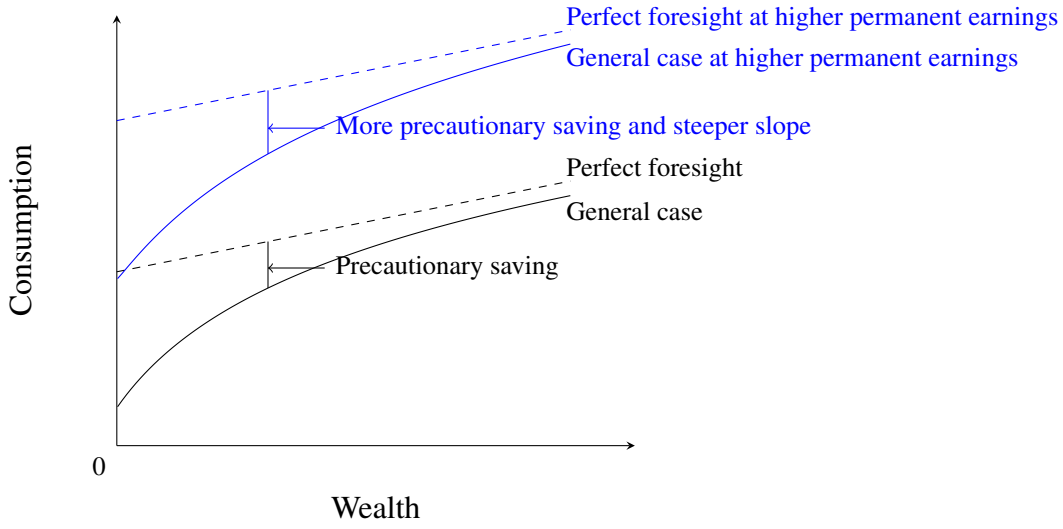


Figure 1: The evolution of consumption with wealth at two levels of permanent earnings

MPC: $\partial MPC_t/\partial e^{p_t} = -(a_t/(e^{p_t})^2)\tilde{c}_2^t(a_t/e^{p_t}, 1, e^{\varepsilon_t}) < 0$.

In the more general case with a non-zero consumption floor $c_0 > 0$, the reasoning presented above can still apply. Indeed, it is not necessary that the MPC remains 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 $c_0 > 0$, the MPC increases when both $a_t > 0$ and e^{p_t} increase proportionally. I also prove in Commault (2024) that c_t is still strictly concave in a_t when $c_0 > 0$. An increase in e^{p_t} is equivalent to a proportional increase in a_t and e^{p_t} (which either keeps the MPC constant or raises it) and a decrease in a_t that puts it back at its initial level (which raises the MPC). Thus, an increase in e^{p_t} raises the MPC in that more general case as well.

Graphical illustration of Proposition 1. Figure 1 presents graphically the mechanism. It plots the evolution of consumption with wealth conditional on other determinants of consumption. The different lines correspond to different levels of permanent earnings (baseline versus high) and different distributions of shocks to future earnings (some shocks versus no shocks). The slope of each line corresponds to the MPC since it measures the partial effect of wealth on consumption.

The dashed black line represents the evolution of consumption with wealth for a baseline level of permanent earnings under perfect foresight, that is, absent any uncertainty about or shocks to future earnings. The evolution of consumption is a straight line in that

case: under perfect foresight people optimally consume a fixed fraction of their lifetime resources—for instance it would be about 3% of the sum of their liquid wealth and discounted flow of future earnings 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. Since people consume a fixed fraction that does not vary with wealth of their lifetime resources, one extra dollar of wealth raises consumption by the same amount at every level of wealth. The MPC, and thus the slope, is constant across the wealth distribution and equal to the value of the fixed fraction. This perfect foresight case is also the 'permanent income' case considered in **Friedman1947**—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 when everything is the same as in the dashed line except that consumers now face uncertainty from transitory and/or permanent 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 wealth reduces the share of risky resources in lifetime resources and reduces precautionary saving. One extra unit of wealth now raises consumption for two reasons: because it raises lifetime expected resources, as in the dashed case, and because it reduces the optimal level of precautionary saving. The response of precautionary saving varies along the wealth distribution. Everything else being equal, at a lower level of wealth, precautionary saving is not only larger but also more sensitive to changes in wealth: one unit increase in wealth relaxes the need to make precautionary saving more at a lower level of wealth, thus raises consumption more at a lower level of wealth. As a result, the MPC decreases with wealth. The slope of the consumption function is steep at low levels of wealth and gets flatter as wealth increases, making consumption concave.

I now plot in blue a situation in which permanent earnings e^{p_t} is higher than in the baseline (black) situation. 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, consumption is simply shifted up: people now consume the same fraction of a larger level of lifetime resources. The dashed blue line is strictly above the dashed black line. Yet the fraction of their lifetime resources that people consume is unchanged. Thus, the MPC is unchanged: the dashed

blue line has the same slope as the dashed black line.

In the presence of earnings uncertainty, people still consume more when they have a higher level of permanent earnings e^{p_t}: the plain blue line lies above the plain black line. However, the precautionary gap between the plain and dashed line is more pronounced in the blue situation: at a given level of wealth, having a higher level of permanent earnings (being in the blue situation instead of the black one) means having a higher ratio of risky to risk-free resources. Consumers with a higher ratio of risky to risk-free resources have a stronger motive to save to accumulate more safe assets: their optimal level of precautionary saving is higher and the gap between the plain and dashed blue lines is larger than the gap between the plain and dashed black lines. At a higher permanent earnings, precautionary saving is not only larger but also more sensitive to changes in wealth: one extra unit of wealth reduces the gap between the plain and dashed lines more when this gap is initially higher. As a result, the slope of the plain blue line is steeper than the slope of the plain black line at the same level of wealth. The MPC is higher at a higher level of permanent earnings, holding wealth constant.

Note that, in the case without a consumption floor, the MPC is homogenous of degree zero in wealth and permanent earnings. This means that increasing permanent earnings (moving from the black to the blue line) and increasing wealth proportionally (moving right on the x-axis by an amount that would be proportional the underlying increase in permanent earnings) would keep the MPC, 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.

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 suggests that permanent earnings is another important determinant. In the case without any consumption floor, the MPC depends on the ratio of wealth to permanent earnings, but is independent of wealth or permanent earnings outside of this ratio. It is thus sufficient to substitute our focus on liquid wealth with a focus on the ratio of liquid wealth to permanent earnings. 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. It is then important to consider both liquid

wealth and permanent earnings separately.

Extension to 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.

I show that Proposition 1 extends to the presence of exogenous borrowing constraints that can be binding, as long as an increase in permanent earnings keeps the current constraint as binding or makes it more binding. When the constraint

can still hold in the presence of an exogenous borrowing constraint, in some situations. In the presence of an exogenous constraint, there are two regimes: one in which the constraint binds, and one in which it does not. When comparing consumers with different levels of permanent earnings who are both in the same regime (both constrained or both unconstrained) and remain so after the one-time unexpected income change, Proposition 1 still holds. When both consumers are constrained, they have the same MPC of one. When both consumers are unconstrained, because the extra saving caused by the possibility of future constraints has the same properties as precautionary saving, Proposition 1 applies—one can go over the analytical proof with exogenous constraints and obtain the same results.

The difficulty is that a higher level of permanent earnings may change the wealth threshold where an exogenous constraint binds, so the consumer with the same wealth but a higher level of permanent earnings might be under a different regime than the other consumer. In the absence of any life-cycle trend in earnings, that is, when the expected value of earnings is the same at all periods ($E_t[y_{t+s}] = cst$ for all $s > 0$), people with a higher level of permanent earnings are less likely to be constrained: an increase in permanent earnings reduces the wealth threshold above which people are no longer constrained.[5] It may move the consumer with a higher permanent earnings outside of the constrained regime. In that case, there exists some levels of wealth at which Proposition 1 no longer holds: a consumer who is constrained can have a higher MPC than a consumer with the same wealth but a

[5]That is because, denoting c_t^* the consumption that would be chosen absent constraints and $\bar{L}_{t+1}$ the exogenous borrowing limit on wealth at the beginning of $t+1$, people are constrained when $a_{t+1} = (1+r)a_t + y_t - c_t^* < \bar{L}_{t+1}$. Because optimal consumption does not increase as much as current earnings with an increase in permanent earnings in the absence of life-cycle trend (see Carroll (2009)), I have $\frac{\partial c_t^*}{\partial e^{p_t}} < \frac{\partial y_t}{\partial e^{p_t}}$. Thus, a marginal increase in permanent earnings reduces the gap between constrained consumption and optimal consumption.

higher permanent earnings, who is not constrained thanks to its higher permanent earnings. However, when there is enough of an upward-sloping trend to earnings, Proposition 1 can be restored even in this knife-edge case.[6] With an upward-sloping trend to earnings, a consumer with higher permanent earnings can be more likely to be constrained, that is, constrained over a wider range of wealth values than a consumer with lower permanent earnings. In that case, Proposition 1 holds. Its interpretation simply changes around the levels where the constraint becomes binding: there can now be two reasons why consumers with a higher permanent earnings respond more to a one-time income shock, because the same exogenous constraint is more binding for them, and because they have a stronger precautionary motive.

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

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 Federal Reserve Bank of New York's Survey of Consumer Expectations (SCE, seeFederal 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 US. A household head is defined as a person in the household who owns, is buying, or rents the home. A household may have multiple co-household heads. Respondents stay on the panel for up to twelve months before rotating out of the panel. The survey started in June 2013. While the Core Survey takes place monthly, its topical modules only take place either every four months or every year. The Labor Survey module, which reports earnings information, take place every four months. The Housing module and the Household Finance survey, which contain information about wealth, take place once a year.[7] In addition to these topical modules Fuster, Kaplan, and Zafar (2020)

[6] When the expected value of future earnings is higher than the expected value of current earnings, it is possible to have $\frac{\partial c_t^*}{\partial e^{p_t}} > \frac{\partial y_t}{\partial e^{p_t}}$.

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

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 **??**.

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

MPCs. I rely on four modules added to the SCE by Fuster, Kaplan, and Zafar (2020). The modules are added 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 makes it possible to avoid a situation in which most people select answer only with 0, 0.5 and 1. 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 one-time unexpected 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 slope of the consumption function is steeper at a higher level of permanent earnings: this means 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. Since the same respondent can be asked several of these questions in the same module thus at the same date, an observation is an MPC question-respondent-date.

The resulting average reported MPC is 0.176. This MPC covers total spending, including spending on nondurable goods, 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 for 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. **OrchardRameyWieland2024** and Borusyak, Jaravel, and Spiess (2022) find the quarterly MPC out of the 2008 US Economic Stimulus Payments to be around 0.25. **ParkerSchildErhardJohnson2021** estimate the quarterly MPC out of the covid stimulus payment to be 0.10 for nondurable spending, with little evidence that durable responded. In a experimental design where participants receive a 300 euros cash transfer, **BoehmFizeJaravel2024** estimate the MPC of total spending out of the transfer to be 0.23 after one month. This is also in line with other surveys asking about the response to an hypothetical shocks. While the earlier survey-based study of Shapiro and Slemrod (2003) asks qualitative questions (would the respondents 'mostly spend' out of an unanticipated gain?), more recent ones ask people to quantify the fraction of an unanticipated gain that they would spend (e.g. **GrazianivanderKlaauwZafar2016** Christelis, Georgarakos, Jappelli, Pistaferri, and Rooij (2019), Bunn, Le Roux, Reinold, and Surico (2018), Parker and Souleles (2019) Crossley, Fisher, Levell, and Low (2021)). The average MPCs are

similar with the one obtained by Fuster, Kaplan, and Zafar (2020).

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 to derive a fading rate. I then apply this fading rate to the following nine months and obtain a yearly MPC.

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 to among fourteen possible bins, ranging from below five hundred dollars to above one million dollars. Its advantage 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 level of savings and investments in accounts other than retirement accounts, the other one is about the share of this hold in saving and checking accounts—thus excluding mutual funds, index funds, stocks and bonds. This measure more precisely captures liquid wealth. However, the survey takes place in August so the response rate among MPC survey respondents is lower.

Demographics. I use demographic characteristics as controls when I examine the effect of permanent earnings on the MPC. Those controls are obtained from the Core module of the SCE. The baseline controls I always include are dummies for the date and the type of MPC question. The set of respondent-specific controls include dummies for the respondent's age group and educational attainment. The set of family-specific controls include dummies for number of household members, for whether there is at least one child (below 18) in the household, and for ranges of total household income.

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 I let those employed face future non-employment risk. The reason for excluding the non-employed is that, to build permanent earnings, I assume that people draw their earnings shocks from the same distributions conditional on the date and fixed 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. I further drop respondents with yearly earnings below $1,885, following Guvenen, Karahan, Ozkan, and Song (2021). To abstain from modeling the education and retirement decision, I also select out people below age 25 and above age

55. Finally, I winsorize the top and bottom 1% of the expected future earnings, earnings, consumption, and variance of future earnings variables. I obtain a final sample of 2,733 observations where I observe jointly the expected future earnings, wealth (the most general variable, non-housing wealth), and the MPC. These observations are coming from 905 respondents. Analyses are then clustered at the repondent level. I deflate all the $ variables using the monthly non-seasonally adjusted Consumer Price Index (CPI), and express them all in March 2016$.

Tables. 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 $e^{p^i_{t+1}} = e^{p^i_t} e^{\eta^i_{t+1}}$ that evolves as a multiplicative random walk and of an innovation x^i_{t+1}. Note that this specification still means that, at a higher level of permanent earnings, the same innovation x^i_{t+1} has a larger impact on y^i_{t+1}. Expected annual earnings is

$$y^i_{t+1} = e^{p^i_t} e^{\eta^i_{t+1}} x^i_{t+1} \;\; \Rightarrow \;\; E^i_t[y^i_{t+1}] = e^{p^i_t} E^i_t[e^{\eta^i_{t+1}} x^i_{t+1}] \tag{3.1}$$

Because the term $E^i_t[x^i_{t+1}]$ is an expectation of an innovation, I assume that it is either the same for all respondents, or the same for all respondents with the same observed characteristics at t. I can filter it out with demographic controls to obtain $e^{p^i_t}$, which is a component that will multiply all future realizations of earnings. That is the intuition for the identification.

A generalized transitory-permanent earnings process. I show that this insight applies the generalized transitory-permanent earnings process proposed in Guvenen, Karahan, Ozkan, and Song (2021). 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[8]

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

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

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

This expression states that the annual earnings of worker i at t, y_t^i, are the product of a dummy for employment status ν_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 it persistence. In practice, this component is virtually permanent because Guvenen, Karahan, Ozkan, and Song (2021) estimate its persistence to be $\rho = 0.991$ and the length of the working age period is less forty years, not infinite. Because I do not seek to estimate this process but simply to find an empirical counterpart to the permanent component, I do not need to assume any specific distributions for these innovations.

The non-employment dummy at t, ν_t^i, is a one/zero dummy for whether the individual is mostly employed or mostly non-employed over the year.[9] I directly observe the probability to be non-employed at t, $p_{\nu_{t-1}^i}$, from a survey question.[10] Thus, I do not need to make assumptions about its particular specification.

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

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

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

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

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 $e^{p_t^i}$ (virtually permanent since $\rho = 0.991$ and lifetime is finite) 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 one unit change in my definition of permanent earnings corresponds to a one dollar increase in annual earnings, at the average age trend and average realization of the current transitory innovation. I denote $perm_t^i$ the re-scaled permanent component

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

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

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

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

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 $ln(E_t^i[e^{\varepsilon_{t+1}^i}])$ are captured by a linear regression over such dummies.[11] Also, ρ is estimated by Guvenen, Karahan, Ozkan, and Song (2021) to be close to one.[12] As a result, the residual res_t^i from a regression of $ln(E_t^i[y_{t+1}]/(1-p_{vt}^i))$ on individual characteristics and year dummies coincides with $p_t^i + \tilde{\alpha}^i$.[13] In the baseline, the individual characteristics z^i that I control for are the gender of the respondent and their willingness to take risks. Indeed, both are likely to affect the distributions from which

[11] This 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 $ln(E_t^i[e^{\eta_{t+1}^i}]) \approx E_t^i[\eta_{t+1}^i] + Var_t^i(\eta_{t+1}^i)/2$ and $ln(E_t^i[e^{\varepsilon_{t+1}^i}]) \approx E_t^i[\varepsilon_{t+1}^i] + Var_t^i(\varepsilon_{t+1}^i)/2$.

[12] Rozsypal and Schlafmann (2023) find that people are partly overestimating the persistence ρ of their earnings. I use self-reported expectations, so the ρ I should use is the one that people believe. Since I set ρ equal to one, this overestimation of the persistence is in line with my specification.

[13] Indeed, it coincides with $p_t^i - \overline{p} + \tilde{\alpha}^i - \overline{\tilde{\alpha}}$ and I set the average sample value of $\overline{p} + \overline{\tilde{\alpha}}$ to zero without loss of generality (any non-zero constant can be captured in $e^{\zeta \bar{z}}$).

people draw shocks. In robustness analyses, I examine the effect of using other fixed characteristics including education and the type of job (private vs public vs NGO).

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 $\overline{ln(y)|_{v_t^i=0}} = \overline{\varepsilon} + \zeta\overline{z} + \overline{g}$. I then take the exponential

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

Since I observe earnings over a longer period than the MPCs, I detrend expected future earnings and build permanent earnings with data from March 2015-March 2019. 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. Furthermore, I note that when expected future employment is observed, my method generalizes to a process with non-employment shocks. Other methods that rely on subjective expectation data include **ArellanoAttanasioCrossmanSancibrian2024** and **ArellanoAttanasioBorellaDeNardiPazPardo2024**. 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 cdf of the shock distribution of permanent earnings, but I identify the value of permanent earnings for each individual: what is the value of $e^{p_t^i}$ for each observation. Other methods to proxy the permanent component of earnings include using simply current earnings and 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 that can be used if a long enough time series is available.

The effect of heterogeneous trends. Guvenen, Karahan, Ozkan, and Song (2021) identify two specifications that fit the data well, which are numbered (5) and (6) in their paper. Here, I draw from their specification (5), which they mention is the simpler one with a reasonable fit, but I can discuss what specification (6) implies. Specification (6) additionally includes an heterogeneous profile $e^{\beta^i t}$ in the realization of annual earnings at t. The

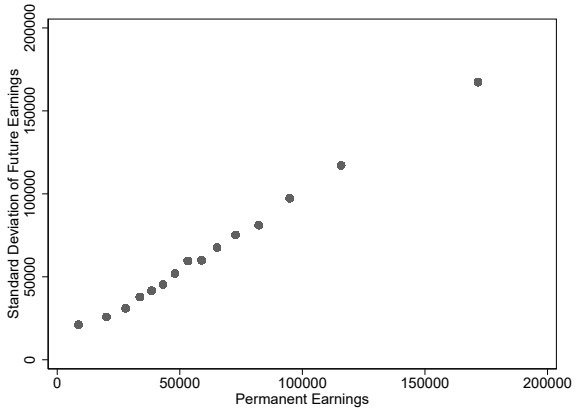


Figure 2: The Standard Deviation of Future Earning And Permanent Earnings

term β^i is distributed with mean zero. In that case, what I measure with my method is $perm_t^{i\prime} = e^{\tilde{\alpha}^i + \zeta\bar{z} + \beta^i(t+1)} e^{p_t^i} e^{\bar{\varepsilon}} e^{\bar{g}} = perm_t^i e^{\beta^i(t+1)}$. This term also multiplies all future realizations of earnings. One difference with the case without trend is that a future realization at $t+s$ with not be multiplied only by $perm_t^{i\prime}$ times $e^{\beta^i(s-1)}$. A larger permanent component of earnings at t coming from a larger β^i should raise the risky component of people's resources—and thus their precautionary motive—even more, because it multiplies future risky earnings y_{t+s} even more than it does current realized earnings y_t.

Permanent earnings and the standard deviation of future earnings. As a first check of my theoretical mechanism, I examine how the standard deviation of future earnings varies with permanent earnings. One typical difficulty with measuring risk is that the ex-ante risk that an individual faces is not observed, only the one outcome that realizes ex-post is. 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 probability they assign to events that would change their annual earnings by given amounts. The questions that I use are about the probability to receive job offers, the levels of these job offers, the probability to accept these job offers, the probability that the employer matches these job offers, and the probability to become unemployed or to leave the labor marker. I use these questions to obtain the different possible states of the world of each worker, with the annual earnings that he or she would have in each of these states of the world and the probability that each of these states of the world realizes.

I can then build the variance of future annual earnings of this worker as the sum of squares of this worker's possible levels of future annual earnings weighted by the probabilities that the worker assigns to each level of future annual earnings. 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. I detail the survey questions and the method that I use to build this variance in Appendix **??**.

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 that scales earnings realizations 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 specification (5.1), 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*.[14] Figure 2 shows that the broadly linear relation holds even between permanent earnings and the *unconditional* standard deviation of future earnings. This means that, although a higher level permanent earnings protects from unemployment in the data, thus reduces future unemployment risk, this effect is not large enough to break the mechanical positive and proportional relation between permanent earnings and the standard deviation 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 the specification (5.1), the coefficient of variation is $CV_t^i = sd_t^i(e^{\eta_{t+1}^i e^{\varepsilon_{t+1}^i - \bar{\varepsilon}}})/E_t^i[e^{\eta_{t+1}^i e^{\varepsilon_{t+1}^i - \bar{\varepsilon}}}]$. 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

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

employed respondents. 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 happens mostly around zero current earnings and is thus essentially driven by people with low attachment to the labor market. I describe those results this in Appendix **??**.

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 $perm_t^i \times e^{\varepsilon_{t+1}-\overline{\varepsilon}}$. 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} was anticipated, it would be present both in my measure of permanent earnings and in the realized innovation to log-earnings, and 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 **??**.

What people are asked about in the SCE is their annual earnings they expect four months from now. The fact that the question is about annual earnings at a short horizon, four months from now, is not per se a problem. Indeed, what I aim to measure is the level of permanent earnings that people have at a given point in time, not the change in their permanent earnings they experience over a period of time. What could be a problem would be seeing no difference between current earnings and expected future earnings because of the short horizon. Yet, people do not report the same values for current annual earnings and expected future annual earnings in March of next year (see Table **??** in Appendix **??**), so the horizon is long enough for them to expect some change. In addition, I verify that current and future expected annual earnings affect the MPC differently (see Table 3 in section 5), which means that the expected annual earnings that people report is not just a resized version of the current earnings they report.

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:

$$MPC_t^i = a_1 + a_2\ perm\ cat_t^i + a_3\ nonhousing\ cat_t^i + a_4 X_t^i + \xi_t^i \tag{4.1}$$

$$MPC_t^i = b_1 + b_2\ perm_t^i + b_3\ (perm_t^i)^2 + b_4\ checking_t^i + b_5\ (checking_t^i)^2 + b_6 X_t^i + \xi_t^i \tag{4.2}$$

The term MPC_t^i is the reported MPC out of a hypothetical shock of respondent i at period t. The term $perm\ cat_t^i$ is a vector of dummies that correspond to the quartile of the permanent earnings of the respondent. The reference category is the first quartile of permanent earnings. Permanent earnings is built as described above. The term $nonhousing\ cat_t^i$ is a vector of dummies that correspond to the quartile of non-housing wealth of the household's respondent. These quartiles are obtained from merging a categorical measure of non-housing household wealth into four new categories that correspond to the quartile of the distribution in my sample. The reference category is the first quartile of non-housing wealth.[15] The term X_t^i designates a vector of control variables. 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. [16] The respondent controls are dummies for the age group, region of residence, and education level of the respondent. The family controls are dummies for the number of household members, number of children, and household income categories. These dummies are described in the data section.

In the second specification (4.2), the term $perm_t^i$ measures the 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 a continuous variable. In this specification, the variables $perm_t^i$ and $checking_t^i$ are standardized.

[15] 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'.

[16] The reference category is \$500 gain today, asked in March 2016. 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.

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 obtained from the estimation of specification (4.1). The first three lines in this column shows that, holding non-housing wealth fixed, moving from the first to higher quartiles of the permanent earnings distribution significantly raises the reported MPC out of a one-time income shock of the respondent. This is in line with the theoretical prediction that I obtain in the first section. 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, and 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 average MPC in this sample is 0.176, so an increase by 0.05 corresponds to a 30% increase for someone at the average MPC. The effect of moving upward in the permanent earnings distribution appears 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 higher quartiles of the non-housing wealth distribution reduces the MPC, consistent with the existing theoretical and empirical literature that established liquid wealth as one of the only non-demographic significant determinant of MPCs. The effect of non-housing wealth is of a broadly similar magnitude as that of moving along the permanent earnings distribution: the point estimates associated with moving to higher quartiles of wealth are between -0.01 and -0.06, while the point estimates associated with the quartiles of permanent earnings are between 0.03 and 0.05.

The second column presents the results when additionally controlling for respondent and household characteristics. The effect of moving from the first quartile of the permanent earnings distribution to higher quartiles is still 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 becomes smaller.

The third column presents the estimation results among the subsample of singles, to whom the standard life-cycle model with one consumer most exactly applies. Indeed, in this subsample, the permanent earnings of the respondent coincides with the permanent

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
MPC							
Permanent earnings Q2	0.052**	0.060***	-0.008				
	(0.021)	(0.022)	(0.047)				
Permanent earnings Q3	0.045**	0.063***	0.053				
	(0.021)	(0.023)	(0.049)				
Permanent earnings Q4	0.028	0.077***	0.092*				
	(0.022)	(0.028)	(0.051)				
Permanent earnings (s.d.)				0.094**			0.083**
				(0.043)			(0.041)
Permanent earnings (s.d.)^2				-0.010			-0.009
				(0.008)			(0.007)
Non-housing wealth Q2	-0.011	0.002	0.009				
	(0.021)	(0.022)	(0.053)				
Non-housing wealth Q3	-0.046**	-0.029	-0.084*				
	(0.021)	(0.022)	(0.048)				
Non-housing wealth Q4	-0.064***	-0.036	-0.109**				
	(0.022)	(0.024)	(0.049)				
Checking and saving accounts (s.d.)				-0.176**	-0.117	-0.196***	
				(0.075)	(0.073)	(0.073)	
Checking and saving accounts (s.d.)^2				0.022	0.010	0.026	
				(0.017)	(0.017)	(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***	-0.117	-0.010	0.026	0.007	-0.020	0.147
	(0.021)	(0.071)	(0.130)	(0.224)	(0.232)	(0.225)	(0.222)
Average MPC	0.176	0.176	0.185	0.218	0.218	0.218	0.218
Question × date controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Respondent controls	No	Yes	Yes	Yes	Yes	Yes	Yes
Family controls	No	Yes	N/A	N/A	N/A	N/A	N/A
Singles only	No	No	Yes	Yes	Yes	Yes	Yes
Observations	2733	2733	493	250	252	252	250
R2	0.130	0.149	0.222	0.232	0.206	0.228	0.213

Standard errors clustered at the respondent level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1: The effect of permanent earnings on the MPC

earnings of the household, not with half of it, and there is only one decision maker in the household. I no longer control for family characteristics since all households are now single-person households. In this smaller sample, only the impact of moving from the first to the fourth quartiles of the permanent earnings distribution is still significant, and it is only so at the 10% level.

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. The permanent earnings and wealth variables are standardized. Either because the continuous specification is less demanding or because the wealth variable is more specifically liquid, the coefficients are more precisely measured. A one standard deviation of permanent earnings ($46,217) around its mean raises the MPC by 0.094. The effect is statistically significant at the 5% level. It is large, as it corresponds to an increase of half the average MPC in the sample. The effect of the square of permanent earnings is smaller and not significant—recall that the level effect already captures the second cross-partial derivative with respect to consumption $\partial MPC_t/\partial e^{p_t} = \partial^2 c_t/(\partial a_t \partial e^{p_t})$, so square captures a third cross partial derivative with respect to consumption. A 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. The effect of the squared value of liquid wealth is small and not significant. The fact that the effect of liquid wealth is now larger than that of permanent earnings likely reflects the fact that the measure of liquid wealth more precisely captures truly liquid wealth. Both results are consistent with the theoretical prediction and points to large and significant effects of both permanent earnings and liquid wealth on the MPC.

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, resulting in a modest estimated effect of liquid wealth on the MPC in the literature: comparing people with almost no liquid wealth to people with a substantial level of liquid wealth, the MPC does not move from one to zero but decreases it much less. The positive effect of the permanent component of earnings on the MPC can account for this fact: people with higher levels of liquid wealth also have higher levels of permanent earnings, so the positive effect of permanent earnings partly offsets the negative effect of liquid wealth.

As a result, the MPC of people with high level of liquid wealth is smaller but not that much smaller. The related prediction is that the estimated effect of liquid wealth should be larger in absolute value when controlling for the offsetting effect of permanent earnings than absent 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 offsetting 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. This has been considered a surprising fact because one would have expected 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 raises the MPC can account for this second stylized fact. While an increase in current earnings increases immediately available resources, as liquid wealth does, an increase in current earnings coming from its permanent component also increases the variance of future earnings which raises the MPC. In contrast, an increase in current earnings coming from its more transitory component raises resources without increasing risk 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 can be small and not significant.

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 not statistically significant. The point estimate of the effect is positive but smaller.[17]

[17]Controlling 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.

Does the ratio of liquid wealth over permanent earnings captures the effect of permanent earnings on the MPC? With isoelastic preferences and a zero consumption floor $c_0 = 0$, so preferences are homothetic, 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 $c_0 > 0$, which generates a form of non-homotheticity, an increase in permanent earnings raises the MPC beyond its effect through the ratio, that is, even when liquid wealth also increases in a way that keeps the ratio of liquid wealth to permanent earnings constant. To test which of the two frameworks fits the data best, I examine the effect of permanent earnings on the MPC when controlling for the ratio of checking and saving accounts amount over permanent earnings rather than controlling for checking and saving accounts amount. The ratio is built with both variables expressed in $ amount and is standardized.

Column (7) presents the results. 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. This is substantially larger than the effect of a one standard deviation increase in liquid wealth alone, consistent with the model prediction. Note that the ratio is not normally distributed but right-skewed, hence large increases in the ratio mostly occur as increases. The second order effect of the ratio is positive and significant. Importantly, beyond this large and positive 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 control for the ratio of wealth over permanent earnings: one needs to control for both wealth and permanent earnings.

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. 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 negative shocks, in particular the effect of moving to a high quartile of permanent earnings. The coefficients are also larger for smaller income shocks, but are no longer significant in the continuous

specification with checking and saving accounts as a measure of liquid wealth, in which the sample is much reduced. The effects are markedly larger and more significant for shocks received today than in the future, typically in three months.

Robustness: bootstrapping. Because the variable of interest, permanent earnings, is built as a residual from a regression, I verify that the results are still significant after accounting for the noise with which this variable is itself measured. One difficulty is that I perform the first step measure of permanent earnings on a different sample than the second-step estimation of (4.1) and (4.2): the first step is with observations from March 2015-2019, while the second-step is on the subset where I observe MPCs, between March 2016 and March 2017. To still examine the standard errors with a bootstrapping procedure, I first modify the estimation to do the first-step estimation over the same sample as the one I use for the second step. This does not meaningfully change neither the point estimates nor their precision. I then move to a bootstrapping-based measure of the standard errors, which accounts for the variability introduced by the first-step regression. The resulting standard errors are smaller than with residuals clustered at the individual level. I present the results in Table **??**, **??**.

Robustness: controlling for optimism. Balleer, Duernecker, Forstner, and Goensch (2021) document an optimistic bias among SCE respondents with a low level of education. The subjective expectations of non-college graduates about their future labor market transitions are on average overoptimistic. In contrast, college graduates have rather precise beliefs. **PfautiSeyrichZinman2024** document that people with lower cognitive skills are overconfident about their future financial situation. Through the mechanism that I document, people with high expected future earnings also have a high MPC. Yet on top of the mechanism that I document, overoptimistic people might have other biases that lead them to consume more out of a one-time income shock. To examine how much of the result is driven by over-optimistic people, I re-run the analyses adding as control a dummy capturing whether people's expected annual earnings t ended up higher than realized or not at $t+1$. This dummy has a non-significant but positive coefficient. It makes little difference in the effect of permanent earnings at t on the reported MPC at t, so it does not seem to be a large driver of my results. These findings are consistent with the facts documented. In particular in **PfautiSeyrichZinman2024**, a first order mechanism is that overoptimistic people consume more than what would be their true optimum, end up with lower wealth, thus higher MPC mostly due to this lower wealth. Thus a small effect of optimism on the

MPC when controlling for wealth can turn into a large unconditional effect of optimism through the dynamic of wealth deccumulation.

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 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 results in **??**.

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 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 $u(.)$ is a log-utility function.

Wealth. People only have access to one perfectly liquid asset in the model. I do a liquid wealth calibration: 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 this the ratio that Kaplan and Violante (2022) use in a paper where they compare different versions of standard incomplete market models. Having only one asset means assuming that people may have illiquid wealth on the side, but that they do not use it to smooth consumption. Kaplan and Violante (2022) document that the distributions of MPCs obtained with such a liquid wealth calibration match the data well, with a fit comparable to that of a two-asset model. With this modeling choice, however, I cannot say anything about total wealth, while

a two-asset model can, 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 \$33,580 (matched +/-1%). The discount factor that matches it is $\beta = 0.973$.

Interest rate. The yearly interest rate on the liquid asset is constant and set to $r = 0.01$, to match the low real interest rate on liquid holdings over the period 2015-2018.[18]

Borrowing limit. In addition to the period budget constraints, people face a borrowing limit on how much they can borrow for consumption purposes. In the baseline calibration, I fix it at a maximum consumption debt of \$5,500 (in \$2014). This is coming from the SCF data about the credit card balance still owned (question x413). The top 90th percentile of unpaid balance is \$5,300 in the 2016 survey (conducted between March and December 2016). This survey wave is the closest to the period I consider, so I set the borrowing constraint at \$5,500.

Lifespan and survival probabilities. People enter the labor market at age 25. They retire at age 62. After retirement, people have a non-zero probability to die at each period from age 62 to age 91. I obtain the survival probabilities from the life tables of the National Center for Health Statistics.[19] 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 that 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 did not have to take a stand on the distribution of the shocks and on the functional form of the probability of non-employment. Here, I follow the distributions and functional forms of Guvenen,

[18] The average 10-Year Real Interest Rate in the US over 2015-2018 is 0.0072 (see Federal Reserve 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.

[19] See `https://www.cdc.gov/nchs/products/life_tables.htm` (accessed May 2024).

Karahan, Ozkan, and Song (2021)

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

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

$$\text{Nonemployment: } \nu_t^i \sim \begin{cases} 0 \text{ (employment) with prob. } 1 - p_{\nu}{}_{t-1}^{i}, \\ 1 \text{ (nonemployment) with prob. } p_{\nu}{}_{t-1}^{i}, \end{cases} \tag{5.3}$$

$$\text{Prob. of nonempl.: } p_\nu(t, e^{p_t^i}) = \frac{e^{\xi_t^i}}{1+e^{\xi_t^i}} \text{ where } \xi_t^i \equiv a_\nu + b_\nu t + c_\nu p_t^i + d_\nu t p_t^i, \tag{5.4}$$

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

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

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

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

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 and social security income. People pay taxes according to the nonlinear tax function of Gouveia and Strauss (1994), $tax(y_t^i) = \tau^b \left(y_t^i - ((y_t^i)^{-\tau^\rho} + \tau^s)^{-1/\tau^\rho}\right)$ parametrized with $\tau^b = 0.258$, $\tau^\rho = 0.768$, $\tau^s = 2.0e-4$ as in Kaplan and Violante (2010).[20]

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

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

MPCs. To compute people's MPCs, I simulate alternative situations where individuals are hit by a one-time unexpected income shock. This is to reproduce the fact that in the survey data, the same respondents are asked about different shocks at the same time. I only include two different shocks instead of 13, for simplicity reasons. 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. In the simulations, the unexpected income gain or loss occur at random, once in the life-time, between age 26 and age 55. I pool together the MPCs out of both shocks. 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).[21]

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

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

[21] The 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.

	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	0.411	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. How close is the model to the data on the earnings, wealth, and MPC dimensions? 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 workers. 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 and commonly used to calibrate this type of models in the recent literature, it is not based on my SCE survey data. Yet the liquid wealth that I observe in the SCE, is not too different. More precisely, I observe two types of liquid wealth. When I convert the non-housing wealth categories 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. This is larger than the level of liquid wealth in the model, but this definition also includes categories of wealth not considered to obtain the SCF ratio. The average amount on checking and saving accounts in the subset of the sample for which it is observed is $26,274. This is about one third of people's non-housing wealth, and close to the average liquid wealth implied by my calibration. The correlation between permanent earnings 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 however lower in the survey data, at 0.328, than in the simulations, at 0.409. Finally, in the simulated data, a little more than one quarter of the people have their liquid wealth at the minimum possible level of -$5,500. 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.

5.3 Specification and main results

Specification. With these simulations data, I estimate a specification close to (4.2), the one that I estimate in survey data when controlling for my more precise measure of liquid wealth, that is, the amount on checking and saving accounts. The controls are the type of MPC question to which the observation corresponds—there are no business cycles so this also corresponds to a Question $\times$ date control—, and age dummies with the same categories as in the empirical estimation. The equation that I estimate is then:

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

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 variations of (4.2). The first column presents the effect of permanent earnings and wealth. The second column presents evidence supporting an explanation for the first stylized fact: the estimated effect of wealth is more modest when failing to control for the effect of permanent earnings. The third column presence evidence supporting an explanation for the second stylized fact: contrary to that of wealth, the effect of current earnings is positive because current earnings is made of both a permanent component, which raises the MPC, and of a transitory component, which reduces it as wealth does. The next three columns present the estimation results of specification (**??**), which is the counterpart to (4.2) for simulated data.

Column (4) shows that, consistent with the theoretical prediction and with the survey results, 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 very close to the increase by 0.220 estimated in survey data. The effect of liquid wealth is negative. The coefficient implies that a one

	Survey data			Simulated data (baseline)		
	(4)	(5)	(6)	(4)	(5)	(6)
Yearly MPC						
Permanent earnings (s.d.)	0.220**			0.182		
	(0.100)					
Permanent earnings (s.d.)^2	-0.024			-0.007		
	(0.018)					
Checking and saving account (s.d.)	-0.410**	-0.273	-0.457***	-0.327	-0.200	-0.300
	(0.175)	(0.169)	(0.169)			
Checking and saving account (s.d.)^2	0.052	0.023	0.062	0.008	0.005	0.007
	(0.040)	(0.040)	(0.039)			
Current earnings (s.d.)			0.174			0.146
			(0.114)			
Current earnings (s.d.)^2			0.015			-0.006
			(0.049)			
Constant	0.060	0.017	-0.048	0.501	0.509	0.539
	(0.522)	(0.541)	(0.524)			
Average yearly MPC	0.509	0.509	0.509	0.512	0.512	0.512
Question × date controls	Yes	Yes	Yes	Yes	Yes	Yes
Respondent controls	Yes	Yes	Yes	Yes	Yes	Yes
Household controls	N/A	N/A	N/A	N/A	N/A	N/A
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

standard deviation increase in liquid wealth reduces the MPC by 0.327. This is similar to, although a little smaller than, the decrease by 0.41 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 to about two-thirds of its value. In the simulations, the coefficient associated with wealth drops from -0.327 to -0.200. Second, the effect of current earnings on the MPC is positive, rather than negative as that of liquid wealth. In the simulations, the coefficient is 0.146. It is lower than that of permanent earnings—equal to about 80% of its value in both simulated and survey data—, 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 very 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 back elements one by one. The simple framework includes a typical transitory-permanent earnings specification. The difference with the richer earnings process is that there is no non-employment shocks, the transitory and permanent innovations are drawn from normal distributions rather than mixture of normals, 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).[22] In this simple framework, I also make social security income proportional to 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. 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 persistent shocks are drawn from a mixture of normal distributions. In Model 5, I additionally reintroduce non-employment shocks. 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 persistent shocks drawn from extreme distributions, the sample is larger since there are no non-employed individuals and less very low-earnings

[22] The variance of the distributions are those of the most likely normal distribution in the Guvenen, Karahan, Ozkan, and Song (2021) 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.

individuals to drop.

The effect of permanent earnings with and without the rich earnings process The first column of the top panel in Table 4 presents the results of specification (4), specification (5), and specification (6) in the most simple model. The results of specification (4) show that permanent earnings has a positive though very small first order effect on the MPC. A one standard deviation increase in permanent earnings at the average level of permanent earnings raises the MPC by 0.0002. This is consistent with the prediction of a theoretical model: absent any extra elements affecting the relation between permanent earnings and the MPC, an increase in permanent earnings raises the MPC. However, with a simple earnings process thus an environment without much risk, this effect remains very small. The effect of liquid wealth is negative. It is also smaller than in the data but less so than the effect of permanent earnings. The results in specification (5) show that, when the mechanism that I identify in the paper 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 mechanism that I identify in the paper is very small, the effect of an increase in current earnings on the MPC is closer to that of an increase in transitory earnings: the coefficient is very small and actually takes a negative sign. The average MPC in this model is very small.

The second column of the top panel presents the results of specification (4), specification (5), and specification (6) in the most simple model plus a progressive social security income scheme. The coefficients all remain very small, and most 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 rather 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 of specification (4), specification (5), and specification (6) in the most simple model to which I add both a progressive social security income scheme and a $5,500 exogenous borrowing limit. This model is close to the one simulated in Kaplan and Violante (2010) or more recently in Kaplan and Violante (2022) for instance. In the estimation of specification (4), the effect of permanent earnings becomes more strongly negative and that of liquid wealth more strongly positive, in

	Model 1 Most simple	Model 2 1+ progressive SSI	Model 3 2+ borrowing constraint
Calibrated β	0.9781	0.9817	0.9709
Specification (4)			
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
Specification (5)			
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
Specification (6)			
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
Specification (4)			
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
Specification (5)			
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
Specification (6)			
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

contrast with empirical estimates. This indicates that the exogenous borrowing constraint alone is not what drives the good fit between the empirical findings and my baseline simulations: introducing it pushes the results further away from the empirical findings. When estimated alone, the effect of liquid wealth is negative but relatively modest (specification (5)). The effect of current earnings on the MPC is modest but negative (specification (6)). The average MPC produced by this model is fairly large, at 0.455. These are consistent with the overall findings of this class of models with a liquid wealth calibration (or two assets) and a relatively tight borrowing constraint: an increase in liquid wealth reduces the MPC but the average effect is relatively modest; this is true although the MPC can be quite large. There is no positive effect of permanent earnings on the MPC, which could explain why the mechanism that I identify, invisible in this class of models, has been overlooked.

The first column of the bottom panel presents the results of specification (4), specification (5), and specification (6) in the most simple model to which I add a progressive social security income scheme, a $5,500 exogenous borrowing limit, and in which the transitory 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 mixture are such that consumers would typically draw from a relatively low variance distribution but there is a small probability to draw from a high variance distribution. 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 small. This model is able to generate qualitatively the two stylized facts that motivate the analysis: the effect of liquid wealth appears 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, when incorporating a mixture of normal distributions, people face substantially larger earnings risk such that 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 of specification (4), specification (5), and specification (6) in the most simple model to which I add a progressive social security income scheme, a $5,500 exogenous borrowing limit, in which the transitory and highly persistent earnings shocks are drawn from a mixture of normal distributions, and in which consumers face non-employment shocks according to the distribution and parameters estimated in Guvenen, Karahan, Ozkan, and Song (2021). The results of specification (4) present coefficients that are very close to the one that I estimate empirically in survey data. The first order effect of permanent earnings on the MPC is 0.24, close to effect of 0.22 in survey data. The first order effect of liquid wealth is -0.35, a little smaller but close to the first order effect of -0.41 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 helps get closer to match 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, at the same level of wealth, consumers with a higher permanent component of earnings have a higher MPC. The mechanism is through precautionary behavior: people with a higher permanent component of earnings have more resources but a larger share of their resources are risky, since they correspond to the expected but yet unrealized flow of future earnings, which is subject to shocks. Everything else being equal, those 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. Everything else being equal, they are more constrained by uncertainty about future earnings: their consumption would increase more than that of others if uncertainty disappeared. A windfall gain that increases the share of certain, realized resources in total expected resources relaxes the need for precautionary saving. People with a higher level of permanent earnings reduces their precautionary saving more and thus consume more out of such a windfall gain.

I find that this theoretical prediction holds true in the New York Fed Survey of Con-

sumer Expectations. First, the distributions of future earnings that people expect are consistent with the existence of a component that multiplies all realizations of future earnings. Indeed, the standard deviation of future earnings increases linearly with expected future earnings, which would be the case where there is permanent component p such that $y = px$: $E[y] = p$ and $sd(y) = psd(x)$. This suggests that, conditional on wealth, the standard deviation of people's resources increases with the permanent component of their earnings. It also means that expected future earnings (conditional on being employed and detrended from the effect of demographic characteristics) can constitute a proxy for the permanent component of earnings. Second, controlling for wealth and demographics, this proxy for the permanent component earnings covaries positively with the respondent's MPCs. A one-standard deviation increase in 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 comparable to that of liquid wealth. More precisely, it is between one-third and two-third as large. Since the two effects go in opposite direction while permanent earnings and liquid wealth are positively correlated, 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. Qauntitatively, I find that, absent any control for permanent earnings, the effect of liquid wealth on the MPC is not significant with a coefficient of -0.12, while with a control for permanent earnings, the effect of liquid wealth becomes statistically significant and the coefficient increases to -0.18. The positive effect of the permanent component earnings on the MPC is also large enough to dominate the negative effect of the transitory component of earnings on the MPC. This can explain why an increase in current earnings—the product of the permanent and transitory component—does not have a negative effect on the MPC although it increases immediate liquidity.

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 MPCs 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 also 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 underes-

timating the effect of liquid wealth, and that current earnings do not have a negative effect on the MPCs. 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

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.

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, Kirill, Jaravel, Xavier, and Spiess, Jann (Apr. 2022). *Revisiting Event Study Designs: Robust and Efficient Estimation*. CEPR Discussion Papers 17247. C.E.P.R. Discussion Papers.

Boutros, Michael (Sept. 2022). "Windfall Income Shocks with Finite Planning Horizons." *Bank of Canada* 2022-40.

Braxton, J. Carter, Herkenhoff, Kyle F, Rothbaum, Jonathan L, and Schmidt, Lawrence (2021). *Changing Income Risk across the US Skill Distribution: Evidence from a Generalized Kalman Filter*. Working Paper 29567. National Bureau of Economic Research.

Broda, Christian and Parker, Jonathan A. (2014). "The Economic Stimulus Payments of 2008 and the aggregate demand for consumption." *Journal of Monetary Economics* 68.S, pp. 20–36.

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.

Commault, Jeanne (2024). "On the Concavity of Consumption in the Permanent Component of Earnings." working paper. see `https://www.jcommault.com/research/`.

Crossley, Thomas F., Fisher, Paul, Levell, Peter, and Low, Hamish (2021). "MPCs in an economic crisis: Spending, saving and private transfers." *Journal of Public Economics Plus* 2, p. 100005. ISSN: 2666-5514.

Fagereng, A., Holm, M., and Natvik, G. (2021). "MPC Heterogeneity and Household Balance Sheets." *American Economic Journal: Macroeconomics* 13.4, pp. 1–54.

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.*

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.

Guvenen, Fatih, Karahan, Fatih, Ozkan, Serdar, and Song, Jae (2021). "What Do Data on Millions of U.S. Workers Reveal About Lifecycle Earnings Dynamics?" *Econometrica* 89.5, pp. 2303–2339.

Guvenen, 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.

Holm, Martin Blomhoff (2018). "Consumption with liquidity constraints: An analytical characterization." *Economics Letters* 167, pp. 40–42. ISSN: 0165-1765.

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, Moll, Benjamin, and Violante, Giovanni L. (2018). "Monetary Policy According to HANK." *American Economic Review* 108.3, 697–743.

Kaplan, Greg and Violante, Giovanni L (2022). *The Marginal Propensity to Consume in Heterogeneous Agent Models*. Working Paper 30013. National Bureau of Economic Research.

Koşar, 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.

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.

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.

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. (2003). "Consumer Response to Tax Rebates." *American Economic Review* 93.1, pp. 381–396.

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.