

How Do Persistent Earnings Affect the Response of Consumption to Transitory Shocks?

Jeanne Commault*

November 2022

Abstract

I show theoretically and empirically that, everything else being equal, people with higher persistent earnings respond more to transitory shocks. Theoretically, people with the same wealth and transitory earnings but higher persistent earnings than others consume more and finance a larger share of their consumption out of their uncertain expected future earnings. Their precautionary motive is thus stronger and their consumption more responsive to transitory shocks. Empirically, in US survey data, an increase in persistent earnings by one standard deviation raises people's consumption response to transitory shocks by 6%-8%. Numerical simulations of a life-cycle model can reproduce these empirical results.

Key words: Consumption Response to Shocks, Heterogeneous Agents, Life-Cycle Models, Household Finance

JEL: D11, D12, D15, E21

*Department of Economics, Sciences Po, 28 Rue des Saints-Pères, 75007 Paris, France; jeanne.commault@sciencespo.fr. I gratefully acknowledges financial support from the Banque de France through its flash grant partnership with Sciences Po.

1 Introduction

How do current earnings affect the way people adjust their consumption when they experience a transitory shock? The answer matters because heterogeneity in the response of consumption to transitory shocks has proved to have large potential aggregate implications.¹ While people have mostly examined how wealth or wealth characteristics (the level and types of assets and debt that consumers hold) influence the response of consumption to transitory shocks, the effect of earnings is potentially important as well. Besides, it can complement our understanding of the effect of wealth, since earnings and wealth correlate. It also has a direct policy relevance, as policy makers typically use income levels to condition the size of fiscal stimulus payments—this was the case to some extent of the 2009 American Recovery and Reinvestment Act and of the 2020 Economic Impact Payments in the US during the covid crisis—, and as Auclert 2019 finds that the (unconditional) covariance between income and consumption responses matters for the design of monetary policy. Finally, understanding the effect of earnings on consumption responses may help interpret some aspects of the Great Recession: a new narrative of this recession gives a more prominent role to prime borrowers with relatively high wages in bearing the wealth loss², but the extent to which these prime borrowers contributed to the consumption plunge that characterized the crisis depends on whether high earners can have large responses to transitory shocks.

The effect of earnings, however, has not been precisely examined theoretically and remains uncertain empirically, with most studies finding non-significant results, a few finding a significant and negative effect, and a few others a significant and positive effect.³

The starting point of this paper is that current earnings are not homogeneous, so one can improve on the theoretical and empirical understanding of their effect by considering their components separately. More precisely, earnings are typically modeled with at least two com-

¹See Kaplan and Violante 2022 for a review on the role played by the heterogeneity of consumption responses to shocks in heterogeneous agent models.

²See e.g. Adelino, Schoar, and Severino 2016, Albanesi, De Giorgi, and Nosal 2017, Foote, Loewenstein, and Willen 2020, Kaplan, Mitman, and Violante 2020.

³Most studies compare the average consumption responses of low, middle and high earners (unconditionally thus without controls) or examine the effect of earnings with only demographics controls. Among them, Parker, Souleles, Johnson, and McClelland 2013, Broda and Parker 2014, Boutros 2021, Fagereng, Holm, and Natvik 2021, Parker, Schild, Erhard, and Johnson 2022 find no significant differences between income categories. Although not significant, the coefficients sometimes reveal a U-shaped pattern in the effect of income, a positive effect in the case of Boutros 2021, and Misra and Surico 2014 find that the median income is higher among people who respond the least or among those who respond most to the 2001 and 2008 stimulus payments. Johnson, Parker, and Souleles 2006 find a significant effect: people in the low income group have a significantly higher consumption response out of the 2001 stimulus payments than those in other groups. Among the studies that control for both demographics and some measure of wealth, Jappelli and Pistaferri 2014 also find the average reported marginal propensity to consume to be significantly higher in the bottom income quintiles. In contrast, two recent papers, Kueng 2018 and Lewis, Melcangi, and Pilossoph 2019, that also examine the effect of income while controlling for a broad range of characteristics including some measure of wealth find that income can have a significant and positive effect on people's consumption response. However, Kueng 2018 examines the response to anticipated income gains (not true shocks), while Lewis, Melcangi, and Pilossoph 2019 finds a positive effect of non-salary income—business and financial income—on the response to the 2008 stimulus payments but not of the rest of income.

ponents, a transitory one and a persistent one. While an increase in the transitory component of earnings just means more cash-in-hand—defined as the sum of risk-free liquid wealth and current earnings—, and should have the same effect as an increase in risk-free liquid wealth, an increase in the persistent component of earnings means more cash-in-hand but also higher future earnings, which are stochastic. It is therefore akin to receiving some non-tradable risky assets with positive but uncertain future dividends, and it might move the response of consumption in a different way than an increase in cash-in-hand.

I make three main contributions: (i) theoretically, I establish that, contrary to a pure increase in cash-in-hand, an increase in persistent earnings raises the consumption response to transitory shocks (for people with positive wealth) in a standard life-cycle model, and I exhibit some proximate conditions sufficient to ensure this holds in a more general set-up; (ii) empirically, I show that people's reported consumption responses are increasing in their persistent earnings when controlling for wealth, other earnings, and demographics: a one standard deviation increase in persistent earnings raises the marginal propensity to consume out of transitory shocks that people report by 0.05, which corresponds to a 6%-8% increase; (iii) numerically, I document that a rich life-cycle model is quantitatively consistent with the empirical results.

I derive the first contribution in a standard life-cycle model with a transitory-persistent earnings process in which the persistent component of earnings simply evolves as a random walk (so its effect on future earnings is in fact permanent). It encompasses as a special case the situation in which persistent earnings constitute an individual, time-invariant level of intrinsic ability. I show that, for somebody with positive wealth, everything else being equal, an increase in this persistent component raises strictly the MPC, that is, the marginal propensity to consume out of cash-in-hand, under some conditions on the utility function that are verified by an isoelastic function (and vice-versa for a decrease in the persistent component).

Intuitively, in this framework, an increase in cash-in-hand leads people to adjust their current consumption for two reasons: first, it raises people's total expected lifetime resources; second, it raises the share of the total expected lifetime resources that is optimal to consume now rather than in the future, because it reduces people's optimal level of precautionary saving, the difference between what they would consume if they faced no uncertainty and what they actually consume. Now, an increase in cash-in-hand raises total resources by the same amount at all levels of persistent earnings, but I show that it reduces precautionary saving more at higher levels of persistent earnings, everything else being equal, for people with positive wealth. The reason why the precautionary motive is enhanced is because, everything else being equal, people with a persistent component of earnings that is twice as large have a level of uncertain future resources that is twice as large but a level of cash-in-hand that is smaller than twice as large when wealth is positive. I then derive proximate conditions that are sufficient for the result to hold in a more general framework that may include other reasons why persistent earnings influence precautionary behavior.

My second contribution is to show that this prediction of the standard model holds true

in US survey data. I consider the more general earnings specification proposed in Guvenen, Karahan, Ozkan, and Song 2021, which among other extensions lets the persistent component of earnings be an AR(1) process, thus not necessarily permanent, and explicitly accounts for unemployment. I show that it is possible to recover an empirical counterpart to this persistent component, which is famously unobserved in survey data, by using expected future earnings. Indeed, future earnings incorporate the persistent component of current earnings but not the transitory component. The method builds on the one in Pistaferri 2001, which uses expectations to identify separately the transitory and permanent shocks that people face. Here I note that, besides identifying the shocks, relying on expectations can also identify the level of the persistent component of earnings, and be adapted to do so in the more general earnings specification I allow for. This strategy expands the set of methods used to identify the persistent component of earnings, which includes using current earnings, using an average of current and past earnings, and, as recently developed in Braxton, Herkenhoff, Rothbaum, and Schmidt 2021, applying a filtering algorithm.

I implement it in the New York Fed Survey of Consumer Expectations (SCE). Using a reduced-form approach, I find that, among employed individuals, my measure of persistent earnings associates significantly and positively with people's reported MPC out of hypothetical transitory earnings shocks (at a one year horizon and for total consumption) when controlling for wealth, other earnings, and demographics. This result is therefore in line with the theoretical prediction of the standard life-cycle model that I expose. Quantitatively, a one standard-deviation increase in persistent earnings associates with a 0.05 level increase in the MPCs out of both negative and positive hypothetical transitory shocks. The average MPC out of a negative shock is 0.796 and the average MPC out of a positive shock is 0.546, which means that, in percentage terms, a 0.05 increase represents a 6% and a 8% increase at the average MPCs. Consistent with the life-cycle model as well, holding persistent earnings constant, the effect of the rest of earnings is negative or not significant. However, when I do not treat persistent earnings and the rest of earnings as distinct, total earnings has no significant effect on the MPCs. These findings confirm the importance of looking separately at the effects of the different components of earnings and explain why studies that do not make this distinction may find overall non-significant results. The results also support the empirical findings of Kueng 2018 and Lewis, Melcangi, and Pilossoph 2019 that MPCs can increase with some earnings components.⁴

Coming back to the reasons why the effect of earnings on the MPCs is important to understand, this result can explain why the literature finds that wealth has only a relatively modest effect on people's MPC (see e.g. Baker 2018, Aydin 2019, Fagereng, Holm, and Natvik 2021,

⁴More precisely, Kueng 2018 considers anticipated income gains and suggests that higher earners might mistakenly respond more than others upon the realization of the gain (rather than when they learn about the shock) because the cost of the mistake is smaller for them. This can make the MPC at the realization of the gain higher for higher earners. My results complement this mechanism by establishing that it is also optimal for higher (persistent) earners to respond more. My results can explain the findings in Lewis, Melcangi, and Pilossoph 2019 if changes in business and financial income are more strongly correlated with changes in future income than other sources of income—which might be the case given the investments it implies.

Ganong, Jones, Noel, Farrell, Greig, and Wheat 2020). Indeed, since persistent earnings and wealth correlate positively but affect the MPC in opposite ways, not controlling for persistent earnings when examining the effect of wealth on the MPC—which is typically the case in the literature since persistent earnings are not observed—leads to underestimating this effect. I confirm this in the survey data. Another implication is that targeting quite narrowly low-earners for fiscal payments does not substantially raise the average MPC out of the payments: although low-earners are likely to have less wealth, they are also likely to have a lower level of persistent earnings. Consistent with this, in the survey data, the average reported MPC is not much higher among people in the 10th earnings percentile than in the whole sample. Finally, a last implication is that even homeowners with high income can be highly responsive to a wealth loss. Consistent with this, in my survey data, the average reported MPC of the homeowners with a level of earnings below the median is the same as the average reported MPC of those with a level of earnings above the median.

My third contribution is to show that a rich life-cycle model that mimics US households is quantitatively consistent with the empirical results. The model encompasses the framework of the theoretical section as a special case. It incorporates the more general earnings specification of Guvenen, Karahan, Ozkan, and Song 2021 that I consider in the empirical section. It also includes a borrowing limit, progressive taxes, transfers, and a retirement period during which people receive social security benefits based on their past earnings and face non-zero death probabilities. I calibrate the model so its average level of wealth matches the average liquid wealth (and not the total wealth) that the individuals in the survey data hold—from the insight of Kaplan and Violante 2014 that wealth has different degrees of liquidity and people use their liquid rather than illiquid wealth to smooth their consumption. In the simulations of this framework, a one standard deviation increase in persistent earnings raises the MPCs out of negative and positive shocks by 0.06 and 0.04, close to the rise by 0.05 in both cases that I estimate in the survey data. Incidentally, the model is able to generate large MPCs, in line with the magnitude of those reported in the survey data.

This rich model I simulate incorporates more channels through which persistent earnings may affect precautionary behavior (in a direction or another) than the standard model of the theoretical section: here, the persistent component of earnings modifies the probability of future unemployment, it is not permanent thus does not multiplies future earnings one-for-one, it influences the taxes, transfers, and retirement income that people pay and receive, and it affects the multiplier on the exogenous borrowing limit. Yet, when I shut down these additional channels, the effect of persistent earnings on the MPCs becomes larger, by approximately one third. This suggests that the standard channel is the main reason for the positive relation I obtain between persistent earnings and the MPCs in these simulations.




Figure 1: The evolution of consumption with wealth at different levels of persistent earnings

2 Persistent earnings and the MPC in life-cycle models

2.1 Graphical intuition

Figure 1 presents graphically the mechanisms that I then prove are at play in a standard life-cycle model with a transitory-persistent earnings specification in which the persistent component of earnings simply evolves as a random walk (so persistent shocks are in fact permanent) and an isoelastic utility. Absent any persistent shocks, the persistent component of earnings corresponds to an individual, time-invariant level of intrinsic ability, which is one possible way to think about it. The plain black line plots the evolution of consumption with accumulated wealth, at a given level of transitory earnings and of persistent earnings. The slope of this function is the MPC, that is, how much consumption responds to a change in wealth. The dashed black line plots this same evolution under perfect foresight, absent any uncertainty about future earnings. The difference between the two corresponds to precautionary saving. This dashed line lies above the plain line: people always consume more (save less) under perfect foresight (because people with an isoelastic utility have a convex marginal utility and therefore make positive precautionary savings).

At a given level of wealth, the MPC (the slope of the plain black line) is steeper than it would be under perfect foresight (the slope of the dashed line): in the presence of uncertainty, a one unit increase in wealth raises lifetime expected resources (the only reason why consumption increases with wealth under perfect foresight), but additionally reduces precautionary saving (the gap between the plain and dashed line decreases as wealth increases). The plain line also increases concavely with wealth: each additional unit of wealth raises lifetime expected resources by the same constant amount (the slope of the dashed line is constant) but reduces precautionary saving less than the previous one (the gap between the plain and dashed line decreases less with wealth at higher levels of wealth).

Now, the effect of precautionary saving is enhanced at higher levels of persistent earnings. I plot in blue a situation in which persistent earnings is higher than in the baseline (black) situation, for instance because the individual has acquired a skill that commands high current and future earnings. At each level of wealth, consumption is higher when persistent earnings are higher (the blue line lies above the black line) but precautionary saving is higher as well (the gap between the dashed and plain lines is larger in blue than in black). The slope of the plain blue line is steeper than the slope of the plain black line because an increase in wealth reduces precautionary saving more when persistent earnings are higher.

2.2 A standard life-cycle model

Consumers' maximization problem Consumers are finite-lived with T the length of their lives. A consumer i chooses consumption expenditures at period t , denoted c_t^i , to maximize lifetime expected utility subject to a number of constraints, as follows:

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

$$\text{with Positive spending: } c > 0, \quad (2.2)$$

$$\text{Budget constraint: } a_t^i = (1+r)a_{t-1}^i + y_t^i - c, \quad (2.3)$$

$$\text{Earnings: } y_t^i = e^{p_t^i} e^{\varepsilon_t^i} e^{\alpha^i} e^{g(t)} \quad (2.4)$$

$$\text{Persistent component: } e^{p_{t+1}^i} = e^{p_t^i} e^{\eta_{t+1}^i}, \quad (2.5)$$

$$\text{Terminal wealth: } a_T^i \geq 0. \quad (2.6)$$

Utility is time-separable and at each period depends only on contemporaneous consumption. The period utility function $u(\cdot)$ is such that marginal utility is positive, decreasing, and convex in consumption: $u'(\cdot) > 0$, $u''(\cdot) < 0$, and $u'''(\cdot) > 0$. The period utility also approaches infinity as consumption approaches zero. The discount factor β captures how much consumers discount utility between two consecutive periods. The positive consumption condition (2.2) imposes that consumption be strictly positive at each period. The budget constraint (2.3) states that consumers earn a stochastic amount y_t^i at each period, and that to store their wealth from one period to another the consumers only have access to a risk-free, liquid asset, with a_t^i the level of this asset at the beginning of period t (or at the end of period $t-1$), that delivers the interest rate r . To simplify the presentation of the mechanism, I assume that $\beta(1+r) = 1$. I discuss later the generalization to $\beta(1+r) \neq 1$ and to β and r time-varying. The earnings process is described by (2.4): earnings are the product of a persistent component $e^{p_t^i}$ and a transitory component $e^{\varepsilon_t^i}$. Earnings may also depend on a fixed effect α^i and on age effect $g(t)$. Incidentally, the transitory-persistent process has initially been applied to the modeling the earnings of individuals (e.g. in Meghir and Pistaferri 2004) but is now used more broadly to model the net income of households (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, and in the empirical and numerical part the transitory-persistent process models earnings. The expression (2.5) states that the persistent component of earnings is a multiplicative random walk process with innovation e^{η^i} . This means that e^{p^i} is not just persistent but permanent, in the sense that its innovation multiplies earnings at each period in the future by the same value it multiplies current earnings without any fading out. In contrast, the realization of e^{ε^i} only affects earnings at t . To obtain that people face a strictly positive amount of uncertainty, I assume $\text{var}(\varepsilon) > 0$. I let the variance of the persistent shocks be possibly equal to zero: $\text{var}(\eta) \geq 0$. When this variance is zero, the persistent component of earnings becomes a fixed characteristic of the individuals (high/low persistent earnings can be interpreted as high/low skill), as for instance in Straub 2019. The persistent and transitory shocks are drawn from distributions that may depend on the consumer's current age and fixed effect. Conditional on those, the shocks are uncorrelated with each other. The terminal condition on wealth (2.6) states that people cannot die with a strictly positive level of debt. The combination of this condition with the budget constraints and positive spending constraints at each period generates a natural borrowing constraint, preventing people from having debt superior to the net present value of the smallest possible realization of their lifetime flow of future earnings. This constraint never binds when marginal utility approaches infinity as consumption approaches zero.

Policy function Conditional on the parameter values and distribution of the shocks, this problem implies that consumption at t is entirely determined by the observation of a_{t-1}^i, e^{p^i} , and e^{ε^i} :

$$c_t^i = c_t^i(a_{t-1}^i, e^{p^i}, e^{\varepsilon^i}).$$

In the remainder of the section, I drop the household index i to ease notations.

2.3 Precautionary behavior and the MPC

Perfect foresight and precautionary saving. Precautionary saving is commonly defined as the 'additional saving that results from the knowledge that the future is uncertain' (Carroll and Kimball (2006)). This characterization is not unique, because there is an infinite number of ways to define the benchmark in which the future is not uncertain.⁵ I use the benchmark that is most common in recent studies: the world in which consumers solve the same problem, with the same current wealth and earnings, except that from period t on their future earnings are

⁵For example, Drèze and Modigliani (1972) study the change in consumption observed when removing uncertainty in a way that keep expected lifetime utility the same, while Kimball (1990a) studies the shift in wealth required for households facing some uncertainty to enjoy the same level of consumption as they would in the absence of uncertainty. Even when one decides to simply set exogenous variables equal to their expected value with probability one, one has to make a choice: some exogenous and uncertain variables are non-linear functions of other exogenous and uncertain variables so they cannot all be equal to their expected value at the same time.

equal to their expected value at t with probability one. I refer to this benchmark as the perfect foresight at t version of the world with uncertainty. I denote with a superscript PF_t the value of a variable under perfect foresight at t .⁶ Precautionary saving, denoted PS_t , is the difference between consumption under perfect foresight at t and consumption:

$$PS_t = c_t^{PF_t} - c_t \quad (2.7)$$

Note that, because the term 'precautionary saving' is now routinely used to refer, not just to the effect of uncertainty, but also to the effect of binding or potentially binding constraints and their interaction with uncertainty, if in more general models people face borrowing constraints other than the natural constraint, I assume them away under perfect foresight so their impact on consumption falls in the category of precautionary saving.

First order condition The first order condition of the problem described by (2.1)-(2.6), also known as the Euler equation, states that an optimizing household equalizes its expected marginal utility over time:

$$u'(c_t) = E_t[u'(c_{t+1})]. \quad (2.8)$$

A household allocates more consumption to the periods associated with a higher expected marginal utility until no unit of consumption can be spent more profitably at another period.

Consumption under perfect foresight. In the absence of uncertainty about future earnings, all the parameters of the maximization problem are known and people can choose in perfect awareness their whole lifetime consumption path. Thus, to equalize their marginal utility over time people equalize their consumption expenditures over time:

$$u'(c_t^{PF}) = u'(c_{t+1}^{PF_t}) \quad (2.9)$$

$$c_t^{PF} = c_{t+1}^{PF_t}. \quad (2.10)$$

There is no consumption growth in this simple model without intertemporal substitution motives ($\beta(1+r) = 1$). Iterating forward, this reasoning implies that $c_{t+s}^{PF_t} = c_t^{PF_t}$ for all $0 < s < T - t$. I substitute each $c_{t+s}^{PF_t}$ with $c_t^{PF_t}$ in the intertemporal budget constraint that emerges from the combination of the period budget constraints with the terminal condition on wealth, and I rearrange. I obtain the classic expression in which consumption is a constant fraction of total

⁶The time t appears in the index PF_t because t denotes the period from which variables are no longer uncertain. For instance, the perfect foresight value at t of wealth at $t+2$ is not the same as its perfect foresight value at $t+1$.




Figure 2: The precautionary premium over different intervals of future consumption

expected resources:

$$c_t^{PF_t} = \frac{1}{l_{t,T}} \underbrace{\left((1+r)a_t + \sum_{s=0}^{T-t} \frac{E_t[y_{t+s}]}{(1+r)^s} \right)}_{\text{Constant fraction of total expected resources}}, \quad (2.11)$$

where the term $l_{t,T} = \sum_{s=0}^{T-t} \frac{1}{(1+r)^s}$ is such that $1/l_{t,T}$ corresponds to the fraction of their total resources that consumers with $T-t$ periods left to live and only intertemporal substitution motives would allocate to period t —e.g. $\frac{1}{l_{t,T}} \rightarrow \frac{r}{1+r}$ when $T \rightarrow \infty$. This fraction is determined by the parameters of the model and independent of the consumers' wealth and earnings.

Consumption in the presence of uncertainty. When the marginal utility function is strictly convex (i.e. $u'''(\cdot) > 0$ which is true of isoelastic utility functions), the presence of uncertainty raises the expected marginal utility of future consumption above the marginal utility of consumption at the expected level of future consumption (from Jensen's inequality $E_t[u'(c_{t+1})] > u'(E_t[c_{t+1}])$). The expected marginal utility of future consumption then rewrites as the marginal utility of expected future consumption minus a premium:⁷

$$u'(c_t) = u'(E_t[c_{t+1}] - \varphi_t) \quad (2.12)$$

$$c_t = E_t[c_{t+1}] - \underbrace{\varphi_t}_{\substack{\text{Premium} \\ > 0}} \quad (2.13)$$

Intuitively, uncertainty raises the marginal utility of future consumption because it opens up the possibility of states of the world in which, because marginal utility is convex, an extra unit of

⁷Kimball 1990b notes that, in a two-period model, the compensating precautionary premium φ_t^* such that $E_t[u'(c_{t+1} + \varphi_t^*)] = u'(E_t[c_{t+1}])$ coincides with the wealth differential that makes people consume the same across different states of the world (with and without earnings uncertainty). What I note here is that, in any multiperiod model, the equivalent precautionary premium φ_t such that $E_t[u'(c_{t+1})] = u'(E_t[c_{t+1}] - \varphi_t)$ coincides with the optimal transfer of consumption between two consecutive periods (it is the amount that equalizes the marginal utility of the certain present and of the the uncertain future).

consumption will be very valuable. As a result, consumers optimally allocate more consumption to the uncertain future than to the certain present, and shift their resources to the future. The graph on the left in Figure 2 illustrates this implication of uncertainty for expected consumption growth. It plots a simple case with a log-utility function and only two possible states of the world at $t + 1$, a low income state and high income state that have the same probability to realize. The terms c_{t+1}^L and c_{t+1}^H denote the (endogenously chosen) values of future consumption in each of these states (in red). The black line plots the marginal utility as a function of consumption. Because this marginal utility is convex and future consumption uncertain (c_{t+1} takes two possible values), the expected marginal utility of future consumption $E_t[u'(c_{t+1})]$ is above the marginal utility of expected future consumption $u'(E_t[c_{t+1}])$: uncertainty raises the expected marginal utility of future consumption, which is driven by the asymmetrically higher marginal utility in the lower consumption state. Because marginal utility is decreasing, current consumption c_t must fall below $E_t[c_{t+1}]$ for its marginal utility $u'(c_t)$ to be as large as $E_t[u'(c_{t+1})]$. The exact amount by which c_t must fall is φ_t such that $u'(c_t) = u'(E_t[c_{t+1}] - \varphi_t)$. Iterating this reasoning forward to future periods, I show in Appendix A.1 that consumption eventually writes as the sum of its perfect foresight value and of a weighted sum of the current and future expected premia φ :

$$c_t = \underbrace{\frac{1}{l_{t,T}} \left((1+r)a_{t-1} + \sum_{s=0}^{T-t} \frac{E_t[y_{t+s}]}{(1+r)^s} \right)}_{\substack{\text{Constant fraction of total} \\ \text{expected resources} \\ = \text{consumption under perfect foresight } c_t^{PF_t}}} - \underbrace{\frac{1}{l_{t,T}} \left(\sum_{j=1}^{T-t} l_{t+j,T} \frac{E_t[\varphi_{t+j-1}]}{(1+r)^j} \right)}_{\substack{\text{Constant fraction of total expected} \\ \text{precautionary growth} \\ = \text{precautionary saving } PS_t}} \quad (2.14)$$

Consumption at t is lower than it would be under perfect foresight at t because the consumers net out from their total expected resources the precautionary consumption growth that they plan to implement. This expected precautionary consumption growth is positive because the terms φ are always positive. They consume a fraction $l_{t,T}$ of what remains instead of a fraction of the full resources. The difference corresponds to precautionary saving. Note that this is not a closed-form expression of consumption but an equilibrium condition, because the variables $\{E_t[\varphi_{t+s-1}]\}_{s=1}^{T-t}$ are not exogenous but jointly determined with c_t .

The two components of the MPC I define the MPC as the effect of a one unit change in the risk-free liquid wealth a_{t-1} on consumption c_t . Indeed, changes in wealth are changes in cash-in-hand and have the same properties as one-time gains or losses in revenue such as stimulus payments or lottery gains: both raise cash-in-hand but do not affect the distribution of future

earnings.⁸ Differentiating both sides of (2.14) with respect to a change in wealth:

$$\frac{\partial c_t}{\partial a_{t-1}} = \underbrace{\frac{(1+r)}{l_{t,T}}}_{=\frac{\partial c_t^{PF_t}}{\partial a_{t-1}}} - \underbrace{\frac{1}{l_{t,T}} \sum_{j=1}^{T-t} \frac{l_{t+j,T}}{(1+r)^j} \frac{\partial E_t[\varphi_{t+j-1}]}{\partial a_t}}_{=\frac{\partial PS_t}{\partial a_t}}. \quad (2.15)$$

Thus, a change in wealth leads people to adjust their current consumption for two reasons: first, it modifies people's total expected lifetime resources, which is the only reason why consumption would respond under perfect foresight; second, it modifies people's optimal level of precautionary saving by shifting the optimal consumption growth between periods, so it modifies the share of the total expected lifetime resources that is optimal to consume now rather than in the future. The extent to which a change in wealth modifies resources is exogenous and does not vary with people's initial wealth or earnings. However, the extent to which it modifies precautionary saving may vary with wealth and earnings. I therefore examine the properties of this latter effect.

Lemma (i) and (ii): wealth and precautionary behavior. In the model described above:

- (i) $\frac{\partial c_t}{\partial a_{t-1}} > \frac{\partial c_t^{PF_t}}{\partial a_{t-1}}$ (so $\frac{\partial PS_t}{\partial a_{t-1}} < 0$), when $\frac{u'''(c)}{-u''(c)}$ strictly decreasing in c ,
- (ii) $\frac{\partial^2 c_t}{\partial a_{t-1}^2} < \frac{\partial^2 c_t^{PF_t}}{\partial a_{t-1}^2} = 0$ (so $\frac{\partial^2 PS_t}{\partial a_{t-1}^2} > 0$), when $\exists k \neq 0$ s.t. $\frac{u'''(c)u'(c)}{(-u''(c))^2} = k$ i.e. $u(c) \in \text{HARA}$,

with HARA denoting the class of hyperbolic absolute risk aversion functions.

Intuition for Lemma (i). The assumption that $\frac{u'''(\cdot)}{-u''(\cdot)}$ is strictly decreasing (which is true of any isoelastic utility function), implies that the convexity of the marginal utility is less pronounced over higher consumption intervals. To see this, note that a small shift $\Delta > 0$ from the interval $[\bar{c}, \underline{c}]$ to the higher interval $[\bar{c} + \Delta, \underline{c} + \Delta]$ is equivalent to staying on the same interval but shifting the marginal utility function from $u'(\cdot)$ to $u'(\cdot) - \Delta \times (-u''(\cdot))$; the concavity of this new function is less pronounced because, when $\frac{u'''(\cdot)}{-u''(\cdot)}$ is strictly decreasing, $-u''(\cdot)$ is a strictly increasing and strictly convex function of $u'(\cdot)$, so removing this relatively more convex element ($-u''(\cdot)$) reduces the convexity of the function. As the convexity lessens over high intervals, so does the need for precautionary saving, and the premium φ_t is smaller over higher consumption

⁸A change in the transitory component of earnings e^{ε_t} also has the same properties as a change in a_{t-1} because both affect consumption only through their impact on cash in hand $(1+r)a_{t-1} + y_t$ and do not affect the distribution of future consumption. However, by construction the sizes of their effect on cash-in-hand are different: a one unit change in a_{t-1} changes cash-in-hand by $(1+r)$ unit, while a one unit change in e^{ε_t} changes cash-in-hand by $e^{p_t} e^{\alpha} e^{g(t)}$. Note that this means that the size of the effect of a change in e^{ε_t} on cash-in-hand depends on e^{p_t} , while the size of the effect of wealth does not. I abstract from this dependency on e^{p_t} by defining the MPC as the partial effect of wealth. If the MPC was defined as the partial effect of e^{ε_t} , this dependency on e^{p_t} would simply create an additional channel through which e^{p_t} affects the MPC.

intervals. The middle graph in Figure 2 plots this result: when the interval $[c_{t+1}^L, c_{t+1}^H]$ shifts upwards by Δ (from the red values to the blue ones), it moves to a region where marginal utility is relatively less convex and the value of φ decreases ($\varphi' < \varphi$).

Now, an increase in wealth does not in general move c_{t+1} upward by the same amount in all states of the world, so it does not simply shift up the distribution of future consumption. However, I prove the result using that, by backward induction, if Lemma (i) is true at $t + 1$, an increase in future wealth raises future consumption at least as much as it would under perfect foresight at $t + 1$, which is the same amount in all states of the world. Therefore, if consumption at $t + 1$ responds as much as it would under perfect foresight at $t + 1$, the premium φ_t decreases and precautionary saving at t decreases. If consumption at $t + 1$ responds strictly more in every state of the world, then consumption at t must respond strictly more than it would otherwise, and precautionary saving at t decreases even more. The detailed proof is in Appendix A.2.

Discussion. This proof of Lemma (i) extends the result of Kimball 1990b that, in a two-period model, at the same level of consumption, the slope of the consumption function is larger in the presence of uncertainty (so the MPC conditional on consumption is larger in the presence of uncertainty) to a multiperiod model and to a comparison of the slopes at the same level of wealth.⁹ The extension from a two-period model to a multiperiod model is not trivial, because two-period models circumvent one important mechanism at play in multiperiod models: the fact that current consumption influences the variance of future consumption—in a two-period model the variance of future consumption simply coincides with the exogenous variance of future income. Previous attempts at extending the result to a multiperiod model had not fully succeeded.¹⁰

Relation to the wealth target narrative. The mechanism I uncover furthers the understanding of Carroll 1997's result that, in a life-cycle model with impatient and infinitely lived people, there exists a level of wealth towards which people would converge if the shocks that realized were such that expected future wealth and realized future wealth coincided, known as the 'wealth target'. The characteristic of this wealth target are therefore that: (1) a decrease in wealth below the target raises saving (thus raises expected wealth growth); (2) an increase in

⁹More precisely Kimball 1990b shows analytically that, in a two-period model, the presence of income risk increases the MPC at a given level of first-period consumption when absolute prudence $u'''(\cdot)/(-u''(\cdot))$ is decreasing. In addition, he notes that, in this two-period model, if consumption is also concave in wealth (which is true when utility is quadratic or, from my next proof, when utility is HARA but not quadratic), then the MPC is also higher in the presence of uncertainty at the same level of wealth. My result establishes this in a multiperiod model, and shows that one does not need consumption to be concave in wealth for the MPC to be higher in the presence of uncertainty at the same level of wealth.

¹⁰Kimball 1990a seeks to extend the results of Kimball 1990b to a model with many periods and with risky assets and suggests that 'the result that the effect of income risk on the marginal propensity to consume out of wealth depends only on whether absolute prudence is increasing or decreasing must be qualified in the presence of either risky securities, or more than two periods'. He provides a more restrictive condition ensuring that the result holds true (and a counterexample but that is not general).

wealth above the target reduces saving (thus reduces expected wealth growth); (3) expected wealth growth is zero at the target; combining the three, expected wealth growth is positive below the target and negative above it.

The result in Lemma (i) shows that (1) and (2) are true at any level of wealth (and even when people are not impatient so no target level of wealth exists): if wealth decreases, the consumption decision problem shifts down, to a region where the convexity of marginal utility is more pronounced and the precautionary motive stronger, so the level of saving increases (and so does expected wealth growth); conversely, if wealth increases, the decision problem shifts up and the level of saving decreases (and so does expected wealth growth). The wealth target defined by Carroll 1997 is the unique level at which this is true *and* expected wealth growth is zero. The conditions of impatience and infinite life simply ensure that a level at which expected wealth growth is zero exists. This means that the behavior of people saving more when their wealth falls and less when their wealth rises cannot be interpreted as their trying to stay close to a given target: this behavior takes place even when a wealth target does not exist, and the reason for this behavior is the changing convexity of the marginal utility function, which is ubiquitous and does not exist only around a given level of wealth.

Intuition for Lemma (ii). When utility is such that $u'''(\cdot)u'(\cdot) = k(-u''(\cdot))^2$ with $k \neq 0$ (which corresponds to a condition of non-quadratic HARA utility, as discussed in the Appendix of Carroll and Kimball 1996, and is true of any isoelastic utility function), the change in the slope of marginal utility $u'''(\cdot)$ is entirely determined by $u'(\cdot)$ and $-u''(\cdot)$: $u'''(\cdot) = k(-u''(\cdot))^2/u'(\cdot)$. This ensures that $u'''(\cdot)$ does not behave as $-u''(\cdot)$ (for $k \neq 0$), and each shift up in the distribution of future consumption reduces the precautionary premium less than the previous shift up (or raises it more than the previous shift if a shift up raises the precautionary premium). The right graph in Figure 2 plots this result: when the interval $[c_{t+1}^L, c_{t+1}^H]$ shifts upwards by Δ for the second time (from the blue values to the orange ones), it moves to a region where the convexity of marginal utility decreases less with a given consumption shift, so this second shift does not reduce the premium as much as the first shift did ($\varphi' - \varphi'' < \varphi - \varphi'$).

Now, again, a change in current wealth does not move c_{t+1} by the same amount in all states of the world. However, by backward induction, if Lemma (ii) holds true at $t + 1$, it is possible to bound the second order effect of current wealth on current consumption. The detailed proof is in Appendix A.3.

Discussion. This proof of Lemma (ii) extends further the result of Carroll and Kimball 1996. Their paper shows that consumption is strictly concave in wealth when utility is HARA with $k > 0$ and $k \neq 1$, for people facing strict earnings uncertainty (their Corollary 1). This proof shows that it actually holds true for any $k \neq 0$, in particular it does for $k = 1$ (exponential utility). The result is consistent with the finding of Toda 2021 that, in life-cycle models with

uncertainty, HARA is a necessary condition for consumption to be (non-strictly) concave.¹¹ This proof also provides a specific intuition: the HARA utility ensures that, following two consecutive shifts up in the distribution of consumption, the change in convexity brought by the second one is larger than the change brought by the first one, either because the decrease in convexity is smaller the second time (when a shift up reduces the convexity of marginal utility) or because the increase in convexity is larger the second time (in the opposite case when a shift up raises the convexity of marginal utility). In both cases, consumption evolves concavely with wealth.

Lemma (iii) and (iv): relation to homogeneity. In the model described above:

$$(iii) \quad c_t = (or \geq) a_{t-1} \frac{\partial c_t}{\partial a_{t-1}} + e^{p_t} \frac{\partial c_t}{\partial e^{p_t}}$$

when $\frac{u'''(c)c}{-u''(c)} > 1$ constant (or strictly increasing/decreasing in c),

$$(iv) \quad 0 = (or \geq) a_{t-1} \frac{\partial^2 c_t}{\partial a_{t-1}^2} + e^{p_t} \frac{\partial^2 c_t}{\partial a_{t-1} \partial e^{p_t}}$$

when $u(c) \in \text{HARA}$ & $\frac{u'''(c)c}{-u''(c)} > 1$ constant (or strictly increasing/decreasing in c).

Intuition for Lemma (iii). By Euler's homogeneous function theorem, the function $c_t = c_t(a_{t-1}, e^{p_t}, e^{\varepsilon_t})$ is homogeneous of degree one in a_{t-1} and e^{p_t} if and only if it is equal to the weighted sum of its derivatives $c_t = a_{t-1} \frac{\partial c_t}{\partial a_{t-1}} + e^{p_t} \frac{\partial c_t}{\partial e^{p_t}}$. Lemma (iii) establishes the conditions under which consumption is homogeneous or compares to this homogeneous case. When relative prudence $\frac{u'''(c)c}{-u''(c)}$ is constant (which implies that utility is isoelastic, in which case relative prudence is strictly larger than one), consumption is exactly homogeneous of degree one in wealth and earnings. This is because, in that case, one can multiply consumption, wealth and persistent earnings by any value λ in the consumers' problem without changing the solution. The assumption of strictly increasing relative prudence ensures that, by backward induction and using the Euler equation, if at $t + 1$ the weighted sum of the consumption derivatives is equal to consumption, then it will be smaller than consumption at t . The assumption of strictly decreasing relative prudence does the opposite. Note that the case with strictly increasing relative prudence encompasses exponential utility functions. The detailed proof is in Appendix A.4.

Discussion. The proof of homogeneity in wealth and in persistent earnings when utility is isoelastic furthers the insight of Carroll 2009 and Carroll 2011 that, in a simple life-cycle model, when utility is isoelastic, it is possible to divide consumption, wealth, and total earnings by current persistent earnings, referred to as a normalization by persistent earnings, and obtain a

¹¹The paper of Gong, Zhong, and Zou 2012 shows however that this HARA requirement can be relaxed in a deterministic consumption problem with $\beta(1+r) > 1$ at each period.

new consumers' problem with one less state variable. While the normalization gives rise to a problem with one less state variable but with a different structure of the shocks, including new shocks to the interest rate and discount factor, the proof of homogeneity tells us that these are irrelevant for the consumption solution: if, conditional on transitory earnings, consumption is a function of wealth and persistent earnings, normalized consumption is *the same function* of normalized wealth and one.¹² The proof of homogeneity also implies that it is equally possible to normalize the consumers' problem by current wealth. Finally, the proof of homogeneity generalizes the result of Straub 2019 that consumption is homogeneous of degree one in current wealth and in persistent earnings in the special case when persistent earnings are a time-invariant individual component not subject to any shock (his Proposition 1).¹³

Intuition for Lemma (iv). By Euler's homogeneous function theorem, the MPC function $(\partial c_t / \partial a_{t-1}) = (\partial c_t / \partial a_{t-1})(a_{t-1}, e^{p_t}, e^{e_t})$, is homogeneous of degree zero in a_{t-1} and e^{p_t} if and only if $0 = a_{t-1} \frac{\partial^2 c_t}{\partial a_{t-1}^2} + e^{p_t} \frac{\partial^2 c_t}{\partial a_{t-1} \partial e^{p_t}}$. Lemma (iv) establishes the conditions under which the MPC, that is, the partial effect of wealth on consumption, is homogeneous or compares to this homogeneous case. When utility is such that consumption is homogeneous of degree one in wealth and persistent earnings hold, the MPC is mechanically homogeneous of degree zero in wealth and persistent earnings. The assumption of strictly increasing relative prudence ensures that, by backward induction and using the Euler equation, if at $t + 1$ the weighted sum of the MPC derivatives is equal to zero, then it will be smaller than zero at t . The assumption of strictly decreasing relative prudence does the opposite. The detailed proof is in Appendix A.5.

¹²I examine this more in details and find that, indeed, the extra shocks cancel out in the expressions that determine consumption because, with an isoelastic utility and a relative risk aversion ρ , the term $(\beta(1+r))^{1/\rho} (1+r)^{-1}$ is the same for the values of β and $(1+r)$ in the original model as in the normalized model. Since this term captures the way the discount factor and interest rate affects consumption, the two affect consumption in the same in the original and normalized models.

¹³The linearity of consumption in the time-invariant individual definition of permanent earnings that Straub 2019 derives from his homogeneity result is not inconsistent with my result because it only holds in the following conditions: if a_{t-1} is proportional to e^{p_t} so that $a_{t-1} = \tilde{a}_{t-1} e^{p_t}$ with \tilde{a}_{t-1} independent of e^{p_t} , then $c_t(a_{t-1}, e^{p_t}) = e^{p_t} c_t(\tilde{a}_{t-1}, 1)$ is linear in e^{p_t} . However, if a_{t-1} is not proportional to e^{p_t} (which happens when people have experienced shocks to their permanent earnings that move them away from proportionality with wealth, or if people have faced no shocks to permanent earnings but their initial assets a_0 were not initially proportional to e^{p_0}), then $c_t(a_{t-1}, e^{p_t}) = e^{p_t} c_t(a_{t-1}/e^{p_t}, 1)$ is not linear in e^{p_t} because the effect of a one unit increase in e^{p_t} on c_t is $c_t(a_{t-1}/e^{p_t}, 1)$, which depends on e^{p_t} .

2.4 Persistent earnings and precautionary behavior

Theorem. In the model described above, for consumers with strictly positive wealth $a_{t-1} > 0$:

- (i) $\frac{\partial c_t}{\partial e^{P_t}} < \frac{\partial c_t^{PF_t}}{\partial e^{P_t}}$ (so $\frac{\partial PS_t}{\partial e^{P_t}} > 0$)
 when $\frac{u'''(c)}{-u''(c)} > 0$ strictly decreasing in c & $\frac{u'''(c)c}{-u''(c)} > 1$ constant or strictly increasing in c ,
- (ii) $\frac{\partial^2 c_t}{\partial a_{t-1} e^{P_t}} > \frac{\partial^2 c_t^{PF_t}}{\partial a_{t-1} e^{P_t}} = 0$ (so $\frac{\partial^2 PS_t}{\partial a_{t-1} e^{P_t}} < 0$)
 when $u(c) \in \text{HARA}$ & $\frac{u'''(c)c}{-u''(c)} > 1$ constant or strictly decreasing in c .

Proof of Theorem (i). Because $\frac{u'''(c)c}{-u''(c)} > 1$ is constant or strictly increasing in c , Lemma (iii) is true and the weighted sum of the consumption derivatives with respect to wealth and persistent earnings is smaller than consumption. This implies that the partial effect of persistent earnings on consumption is:

$$\frac{\partial c_t}{\partial e^{P_t}} \leq \frac{c_t}{e^{P_t}} - \frac{a_{t-1}}{e^{P_t}} \frac{\partial c_t}{\partial a_{t-1}} < \frac{c_t^{PF_t}}{e^{P_t}} - \frac{a_{t-1}}{e^{P_t}} \frac{\partial c_t^{PF_t}}{\partial a_{t-1}} = \frac{\partial c_t^{PF_t}}{\partial e^{P_t}}. \quad (2.16)$$

I move from the second to the third expression using that, when $\frac{u'''(c)}{-u''(c)} > 0$, people make precautionary saving and $c_t < c_t^{PF_t}$ so $\frac{c_t}{e^{P_t}} < \frac{c_t^{PF_t}}{e^{P_t}}$ because, that $a_{t-1} > 0$, and that, when $\frac{u'''(c)}{-u''(c)} > 0$ is strictly decreasing in c , Lemma (i) is true and $\frac{\partial c_t}{\partial a_{t-1}} > \frac{\partial c_t^{PF_t}}{\partial a_{t-1}}$. I move from the third to the fourth expression using that, because under perfect foresight consumption is linear in wealth and persistent earnings it is also homogeneous of degree one in wealth and persistent earnings, to substitute $\frac{c_t^{PF_t}}{e^{P_t}} - \frac{a_{t-1}}{e^{P_t}} \frac{\partial c_t^{PF_t}}{\partial a_{t-1}}$ with $\frac{\partial c_t^{PF_t}}{\partial e^{P_t}}$.

Intuitively, increasing persistent earnings by one unit is equivalent to multiplying both persistent earnings and wealth (by $(1 + 1/e^{P_t})$), and then shifting wealth down (by $-\frac{a_{t-1}}{e^{P_t}}$ units) to account for the fact that wealth was not actually multiplied by $(1 + 1/e^{P_t})$. When consumption is homogeneous of degree one in wealth and persistent earnings, multiplying both variables by $(1 + 1/e^{P_t})$ induces consumption to increase by $\frac{c_t}{e^{P_t}}$ (its value multiplies by $(1 + 1/e^{P_t})$). This term is smaller for people facing uncertainty because they make precautionary saving, so their initial level of consumption is smaller. That is what the first term captures. When wealth is positive, a shift in wealth by $-\frac{a_{t-1}}{e^{P_t}}$ corresponds to a decrease in wealth. The decrease is larger for people facing uncertainty, because for them a decrease in wealth does not only reduce resources but also shifts the distribution of their future consumption to a region where the convexity of marginal utility is more pronounced and raises their optimal level of precautionary saving. This is what the second term captures. This second effect implies that the ratio of precautionary saving over persistent earnings, not just the level of precautionary saving, increases with persistent earnings.

When wealth is initially negative, the first term $\frac{c_t}{e^{P_t}}$ is still larger than it would be under

perfect foresight, but the second one is now smaller than it would be under perfect foresight: people are initially in debt and this debt becomes relatively smaller when their persistent earnings increase. Because people with a precautionary motive are more sensitive to variations in their wealth, this relative debt reduction stimulates their consumption more than it would under perfect foresight.

Proof of Theorem (ii). Because $\frac{u'''(c)c}{-u''(c)} > 1$ is constant or strictly decreasing in c , Lemma (iv) is true, and the weighted sum of the MPC derivatives with respect to wealth and persistent earnings is larger than zero (Lemma (iv)). This implies that the partial effect of persistent earnings on the MPC is:

$$\frac{\partial^2 c_t}{\partial a_{t-1} e^{p_t}} \geq -\frac{a_{t-1}}{e^{p_t}} \frac{\partial^2 c_t}{\partial a_{t-1}^2} > -\frac{a_{t-1}}{e^{p_t}} \frac{\partial^2 c_t^{PF_t}}{\partial a_{t-1}^2} = 0. \quad (2.17)$$

I move from the second to the third expression using that $a_{t-1} > 0$ and that, when $u(c)$ displays HARA, Lemma (ii) is true and $\frac{\partial^2 c_t}{\partial a_{t-1}^2} < \frac{\partial^2 c_t^{PF_t}}{\partial a_{t-1}^2} = 0$.

Intuitively, again, increasing persistent earnings by one unit is equivalent to multiplying both persistent earnings and wealth (by $(1 + 1/e^{p_t})$), and then shifting wealth down (by $-\frac{a_{t-1}}{e^{p_t}}$ units) to account for the fact that wealth did not actually increase. The scaling up does not affect the MPC, but the reduction in wealth raises the MPC when consumption is concave in wealth, because an increase in wealth reduces precautionary saving more at lower levels of wealth—or raises it less in the case with absolute prudence not decreasing. The MPC is therefore higher at higher levels of persistent earnings because an increase in wealth reduces precautionary saving more at higher levels of persistent earnings—or raises it less in the case with absolute prudence not decreasing.

Time-varying demographics, discount factors, and interest rates. In a model with time-varying demographic characteristics z , discount factors β and interest rates r , such that the value function is $V_t(a_{t-1}, e^{p_t}, e^{\varepsilon_t}) = \max_c u(c) + \beta_{t+1} e^{\delta z_{t+1} - z_t} E_t [V_{t+1}(a_t, e^{p_{t+1}}, e^{\varepsilon_{t+1}})]$ and the period budget constraint is $a_t^i = (1 + r_t)a_{t-1} + y_t - c$, the Euler equation becomes:

$$u'(c_t) = E_t[u'(c_{t+1})]R_{t,t+1}, \quad (2.18)$$

with $R_{t,t+s} \equiv \prod_{k=1}^s \beta_{t+k} (1 + r_{t+k-1}) e^{\delta(z_{t+k} - z_{t+k-1})}$ not necessarily equal to one. Redoing the computations in Appendix A.1, the expression of consumption in the presence of uncertainty is preserved but with a different value of $l_{t,T} = \sum_{s=0}^{T-t} \frac{R_{t,t+s}^{1/p}}{\prod_{k=0}^s (1+r_{t+k})}$. The Lemma (ii), (iii) and (iv) are unaffected by the presence of the additional term $R_{t,t+1}$, because I do not need to compare the expressions with their perfect foresight values: in Lemma (ii) I compare the general value to zero, in Lemma (iii) and (iv) I do not use the perfect foresight counterfactual. Thus, Theorem (ii) is unaffected as well. The Lemma (i) still holds when the function $g(\cdot) = (-u'') \circ (u')^{-1}(\cdot)$

is such that for any y , R , $\frac{g'(yR^{-1})R}{g'(y)R} - \frac{g(yR^{-1})}{g(y)} \geq 0$, as in discussed in its proof. The isoelastic utility function and a range of exponential utility functions verify this additional condition. Theorem (i) therefore still holds under this additional condition.

General conditions The mechanisms I show are at play in this standard life-cycle model might hold true more generally, when the persistent component of earnings shifts the precautionary motive not just because it multiplies current and future uncertain earnings (and not current wealth) but also because it affects other parameters or variables of the model. Indeed, the change in precautionary saving and slope described in Figure 1 might equivalently be caused by these other channels and still hold true in more general frameworks.

I thus generalize the model to a situation in which people must decide on their consumption, which they finance out of different resources including one that represents persistent earnings and one that represents risk-free liquid wealth. I do not specify the exact problem that people solve, nor the exact shape of the earnings process they face, except that, absent uncertainty, their consumption is linear in risk-free liquid wealth and overall homogeneous in their different resources. I then exhibit in Appendix A.6 two proximate conditions that are sufficient to ensure that Theorem (i) holds, and two that are sufficient to ensure that Theorem (ii) hold. The conditions that guarantee Theorem (i) are that an increase in the resources unrelated to persistent earnings reduces precautionary saving, and that consumption is homogeneous of degree one in the different resources (or that the weighted sum of the consumption derivatives with respect to each resource is smaller than consumption). The conditions that guarantee Theorem (ii) are that an increase in the resources unrelated to persistent earnings raises the response of precautionary saving to a change in risk-free liquid wealth (e.g. reduces the extent to which precautionary saving decreases with risk-free liquid wealth), that is, reduces the precautionary component of the MPC, and that the MPC is homogeneous of degree zero in the different resources (or that the weighted sum of the MPC derivatives with respect to each resource is larger than zero).

3 Empirical measure of persistent earnings

3.1 Survey of Consumer Expectations

Data. I rely on data from the Survey of Consumer Expectations (SCE) of the Federal Reserve Bank of New-York. It is a monthly online survey with a rotating panel of about 1,300 household heads based in the US—with a household head defined as the person in the household who owns, is buying, or rents the home (a household may have multiple cohousehold 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 every four months or every year. As a result, I only observe income, consumption and wealth simultaneously once every year, around November/December. This means there is

no panel dimension in the analyses that include earnings, consumption and wealth—but there is a panel dimension in some analyses of earnings. I describe the way I match these different modules in detail in B.1. Also, because not all the modules started in 2013, the period over which I observe jointly these variables is November 2015–November 2018.¹⁴

Variables. I obtain current annual earnings, expected future annual earnings in four months, and the probability to be employed in four months from questions in the Labor Market module of the SCE. I also use questions about the probability of future earnings-related events to build an alternative, indirect, measure of expected annual earnings as well as an individual-level measure of the variance of future earnings.

I build the MPCs out of negative and positive transitory shocks from questions in the Household Spending module of the SCE. For the MPC out of a negative shock, the survey asks respondents to consider an hypothetical situation in which their annual income next year would be 10% lower (the survey takes place in December, so next year corresponds to the next twelve months). The respondents first state whether they would cover the loss by cutting their household spending, cutting their household saving, borrowing more, or engage in combinations of the three. Those who would choose combinations are asked to quantify how much of the drop in income would be absorbed by each of these actions. I assume people answer over the same time horizon as the income change, that is, over the next year. I build the MPC out of a positive shock in the same way. I verify in Appendix B.2 that, although the transitory nature of the shock is not strongly stressed, the answers are of the same magnitude as in other surveys with hypothetical MPC questions in which the fact that the shock is transitory is very explicit. Due to the ambiguity of the borrowing and debt repayment answers that are proposed, in my baseline definition of the MPCs, I treat a reported debt repayment as an increase in consumption over the course of the year (and a reported increase in borrowing as a decrease in consumption over the course of the year but this answer is infrequent). Indeed, I argue in Appendix B.2 that people who pay off some of their debt with an income gain are likely to consume more than they would have if they still had the debt to repay, even to pass on all of the debt repayment to consumption and end up with the same debt they initially had if they are constrained, although they state they would use the money for debt repayment. Similarly, although people report they would take on new debt in response to an income loss are rare in the sample, they are likely to consume less than they would if they did not have this new debt. In addition, I show that the average MPCs I obtain with the inclusion of debt repayment and new debt are much more consistent with MPC estimates from natural experiments than without this inclusion. The recent study of Karger and Rajan 2020 also presents some results bundling together spending and debt repayment, as I do here.

I build wealth from a question in the Housing module of the SCE asking respondents to se-

¹⁴See Armantier, Topa, Klaauw, and Zafar 2017) for technical background information on the SCE, and www.newyorkfed.org/microeconomics/sce.html for additional information

lect which of 14 possible categories of non-housing wealth their household belongs to. I check in Appendix ?? that the responses are consistent with other questions about non-housing wealth from the Household Finance module—which I do not use because I would lose many observations. These variables are all deflated with the Consumer Price Index (CPI) and expressed in 2014\$.

I use demographic characteristics in two instances: to net out their effects from individual earnings, and to net out their effects from household consumption. According to the specification that I use, the earnings-related demographic variables are fixed individual characteristics (I use gender, educational attainment, and willingness to take risks dummies to capture them) plus age and period dummies. The consumption-related demographic variables that I use to detrend consumption are the same (excluding gender because I move from the individual to the household level), plus family size and state of residence dummies. I obtain all those variables from the Core module of the SCE.

I present the text of the questions and detail the way I build these variables in Appendix B.2.

Selection. I exclude non-employed individuals. This is because one of the assumptions I make to build persistent earnings is that, conditional on fixed individual characteristics, age dummies, and period dummies, people draw earnings shocks from the same distributions. I find evidence that this holds well among employed respondents, but no longer does when I include non-employed respondents, which suggests that the distributions people draw their shocks from depend on their employment status. I further drop respondents with yearly earnings below \$1,885, following Guvenen, Karahan, Ozkan, and Song 2021. To abstain from modeling the retirement decision, I also select out people above age 55. Finally, I trim the top and bottom 1% of the expected future earnings, earnings, consumption, and individual variance variables (recoding the top and bottom values as non-reported so that the order in which I trim the variables does not matter).¹⁵ I present descriptive statistics of my main variables in Appendix B.3.

3.2 Measuring persistent earnings

A flexible transitory-persistent process. I let annual earnings be a flexible transitory-persistent process, richer than the simple process I use to establish my theoretical point. More precisely, I model the annual earnings y_t^i of individual i at year t with a general specification drawn from the one proposed in Guvenen, Karahan, Ozkan, and Song 2021, which they show fits adminis-

¹⁵I drop the 22 people whose reported responses to the MPC questions (what they would do with the loss or gain) do not add up to 100% so I do not further trim the MPCs.

trative US data on earnings well.¹⁶ It is:

$$\text{Annual earnings: } y_t^i = \underbrace{(1 - v_t^i)}_{\text{Empl. status}} \underbrace{e^{p_t^i}}_{\text{Persistent}} \underbrace{e^{\varepsilon_t^i}}_{\text{Transitory}} \underbrace{e^{\alpha^i}}_{\text{Fixed effect}} \underbrace{e^{g(t)}}_{\text{Age trend}} \quad (3.1)$$

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

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

This expression states that annual earnings are the product of a dummy for employment status v_t^i , a persistent component $e^{p_t^i}$, a transitory component $e^{\varepsilon_t^i}$, a fixed effect e^{α^i} , and a deterministic age trend $e^{g(t)}$.

The term p_t^i , that is, the log of the persistent component is an AR(1) process with $\rho \leq 1$ the AR(1) coefficient. The term η_t^i denotes a persistent shock, which affects people's earnings for the rest of their working life, although the magnitude of its effect decreases over time when $\rho < 1$. The term ε_t^i is a transitory shock that only affects earnings for the current year. The persistent and transitory shocks are drawn independently and are not serially correlated. Their distributions might vary with people's employment status, fixed demographic characteristics, age, and with the period but, conditionally on this, the distributions are the same across households.

The nonemployment dummy at t , v_t^i is a one/zero dummy. Note that this prevents the possibility of nonemployment spells shorter than a year: the reason for this assumption is that, when Guvenen, Karahan, Ozkan, and Song 2021 let the length of the nonemployment spells be a parameter that they estimate, they find it to be very close to one (which corresponds to one year). The probability to be nonemployed at t , $p_{v_{t-1}^i}$ may depend on any characteristic of the individuals at $t - 1$. Contrary to Guvenen, Karahan, Ozkan, and Song 2021, for simplicity, I assume it is entirely determined at $t - 1$, without allowing it to only be determined after the realization of the persistent earnings shock η_t^i . Yet, because I can allow the distribution of persistent earnings shock at t to depend on characteristics at $t - 1$, a correlation between η_t^i and $p_{v_{t-1}^i}$ can still exist in my specification.

Note that this generalized process lets persistent earnings affect precautionary behavior through more channels than the earnings process assumed in the standard life-cycle model of the theoretical section. In particular, here, persistent earnings may affect the probability to be nonemployed next period, therefore affect the variance of future earnings. In fact, Guvenen, Karahan, Ozkan, and Song 2021 show via numerical simulations that simply letting the probability to be nonemployed depend on persistent earnings makes it possible to reproduce the stylized fact that the standard deviation of log-earnings growth decreases with earnings. Also,

¹⁶Guvenen, 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).

because the persistence ρ may be smaller than one, an increase in current persistent earnings does not multiply future earnings exactly by the same amount, which might alter its effect on precautionary behavior.

Rescaling to ease the interpretation of changes in persistent earnings. One difficulty is that the effect of a one unit change in e^p on earnings is unclear, because the relative magnitude of e^p compared to that of the other components is unclear. To simplify the interpretation, I rewrite the log of the transitory component, fixed effect, and age trend in terms of deviations around their means (denoted with an overbar). I then incorporate the exponential of their means, which are constant across households, into a normalized persistent component, denoted $perst_t^i$

$$y_t^i = \underbrace{(1 - v_t^i)}_{\text{Empl. status}} \underbrace{e^{p_t^i} e^{\bar{\varepsilon}} e^{\bar{\alpha}} e^{\bar{g}}}_{\text{Resc. persistent} \equiv perst_t^i} \underbrace{e^{\varepsilon_t^i - \bar{\varepsilon}}}_{\text{Resc. transitory}} \underbrace{e^{\alpha^i - \bar{\alpha}}}_{\text{Resc. fixed effect}} \underbrace{e^{g(t) - \bar{g}(t)}}_{\text{Resc. age trend}} \quad (3.4)$$

A one unit increase in rescaled persistent earnings $perst_t^i$ then coincides with a one dollar increase in earnings for an employed individual at the average sample values of ε , α , and $g(t)$. Note that, because the exponential function is convex, from Jensen's inequality, the average sample values of the rescaled transitory component, fixed effect, and age trend are above one. Thus, by construction, among employed individuals, average rescaled persistent earnings are smaller than average earnings.¹⁷

Using detrended expected future earnings as a measure of persistent earnings. I use the fact that expected future annual earnings at $t + 1$ depend on current persistent earnings but not on current transitory earnings to identify the persistent component of current earnings. More precisely, expected future earnings depend on $e^{p_t^i}$, on the fixed effect component e^{α^i} , which I assume can be completely captured by fixed, time-invariant demographics, on the expected values of the future shocks $E_t^i[e^{\eta_{t+1}^i}]E_t^i[e^{\varepsilon_{t+1}^i}]$, which are constant across households conditional on employment status, fixed demographics, age, and the period, on the deterministic age trend, and on the probability (determined at t) to be employed at $t + 1$

$$E_t^i[y_{t+1}^i] = (e^{p_t^i})^\rho \underbrace{e^{\alpha^i}}_{\text{Constant cond. on dem.}} \underbrace{E_t^i[e^{\eta_{t+1}^i}]E_t^i[e^{\varepsilon_{t+1}^i}]}_{\text{Constant cond. on period \& fixed dem.}} \underbrace{e^{g(t+1)}}_{\text{Age trend}} (1 - p_{v_t^i}). \quad (3.5)$$

Dividing expected future annual earnings by the probability to be employed at the next period

¹⁷Formally, the average rescaled transitory component in the population is $E[e^{\varepsilon^i - \bar{\varepsilon}}] > e^{E[\varepsilon^i] - \bar{\varepsilon}} = 1$, the average rescaled fixed effect is $E[e^{\alpha^i - \bar{\alpha}}] > e^{E[\alpha^i] - \bar{\alpha}} = 1$, and the average rescaled age trend is $E[e^{g(t) - \bar{g}(t)}] > e^{E[g(t)] - \bar{g}(t)} = 1$. The average annual earnings among employed people is then the average rescaled persistent earnings, multiplied by three terms strictly larger than one. As a result, average annual earnings are always strictly larger than average rescaled persistent earnings.

and taking the log of the resulting term yields

$$\underbrace{\ln\left(\frac{E_t^i[y_{t+1}^i]}{(1-p_{v_t}^i)}\right)}_{\text{Observed}} = \underbrace{\rho}_{=0.991 \text{ from GKOS}} p_t^i + \underbrace{\alpha^i + \ln(E_t^i[e^{\eta_{t+1}^i}]E_t^i[e^{\varepsilon_{t+1}^i}]) + g(t+1)}_{\text{Captured through fixed demographics, age, and period dummies}}. \quad (3.6)$$

Thus, among employed people, the residual from a regression of $\ln(E_t^i[y_{t+1}^i]/(1-p_{v_t}^i))$ on demographic dummy variables capturing the fixed effect component of income (dummies for gender, educational attainment, and willingness to take risks) and on dummy variables for the age category and the period should coincide with ρp_t^i .¹⁸ I denote res_t^i this residual. I divide it by $\rho = 0.991$, which is the parameter value estimated in Guvenen, Karahan, Ozkan, and Song 2021 (denoted GKOS in (3.6))—for their specification (5), the one I use. This gives me a measure of p_t^i . To rescale, I multiply this term by the average log-income among employed respondents, denoted $\overline{\ln(y)}|_{empl}$. Finally, I take the exponential. This gives me a measure of the rescaled persistent earnings $perst_t^i = e^{p_t^i} e^{\bar{\varepsilon}} e^{\bar{\alpha}} e^{\bar{g}(t)}$:

$$e^{\frac{1}{\rho} res_t^i \times \overline{\ln(y)}|_{empl}} = e^{p_t^i} e^{\bar{\varepsilon}} e^{\bar{\alpha}} e^{\bar{g}(t)} = perst_t^i.$$

Note that the dummy for the willingness to take risk, which is one of the fixed effect demographics I use, helps control for heterogeneity in risk-aversion but also for some heterogeneity in optimism—if people who declare themselves to be more willing to take risks are in part overoptimistic.¹⁹

Summary statistics. Table 1 presents summary statistics on the variables I use to build persistent earnings and on my resulting measure of persistent earnings, among the respondents in the final sample—those in the selected sample for whom I jointly observe my measure of persistent earnings, my categorical measure of wealth, at least one of the two MPCs, and the household consumption demographics (state of residence and family size) that I use in the main specification. The values are in 2014\$. The first two lines show that people do expect some change in their annual earnings four months from now: annual earnings in November are lower than

¹⁸It would not be feasible to run a fixed-effect regression in this context—although I observe earnings and demographic variables (the variables required for this regression) every four months. Indeed, using it to select out the fixed effect and treating the residual as my measure of persistent earnings would mean I would get rid of the persistent component at the beginning of period t and only capture the changes in persistent earnings that take place during the year. These variations are too small to significantly covary with the MPC (theoretically these variations are zero if shocks are yearly). Using the fixed effect regression and treating the fixed effect component as my measure of persistent earnings does not eliminate the fixed demographics that also covary with the MPC in possibly different directions than persistent earnings and blur the relation I seek to identify.

¹⁹Also on this, Balleer, Duernecker, Forstner, and Goensch 2021 document an optimistic bias, in the data from the SCE, that households have when they form subjective expectations about their future labor market transitions. However, the paper also documents that this bias strongly correlates with educational achievement—college graduates having rather precise beliefs and non-college graduates much less so—, so the education dummies that I use should capture it.

	Mean	Coef. of var.	Obs.
Annual earnings	63,471	0.637	1,117
Expected annual earn.	64,655	0.636	1,117
Expected proba to be empl.	0.976	0.072	1,117
Expected annual earn. cond. on empl.	66,392	0.636	1,117
Persistent earnings.	59,687	0.551	1,117

Table 1: Descriptive statistics on my measure of persistent earnings

expected annual earnings four months later (\$63,471 vs \$64,655 on average) and as volatile (the coefficient of variation is 0.637 vs 0.636). The third line shows that employed people put a very high probability on their still being employed in four months (0.972). The average value of expected earnings conditional on employment, built by dividing expected annual earnings over the probability to be employed, increases to \$66,392. Using this expected annual earnings conditional on employment to extract the persistent component of current earnings, as detailed above, the average rescaled persistent earnings is \$59,687 and the coefficient of variation is smaller (0.551). This suggests that a substantial part of the variation of expected annual earnings in the sample is coming from the demographics and period dummies—capturing differences in the fixed effect component and in the distribution of future shocks.

Robustness. Importantly, it is not particularly problematic that what I observe is expected annual earnings four months from now rather than a year from now. This is because what I seek to measure is the level of persistent earnings that people have at a given point in time, rather than its variations, so the expectations horizon does not matter as long as expected future earnings do not include any transitory component—which I check in the next section—but does include a persistent component. In addition, I explain in Appendix B.4 that the method does not rely on any specific assumption about the timing of the persistent and transitory shocks within the year. I also explain in Appendix B.4 that the method can also accommodate different assumptions about what people have in mind when they report their expected annual earnings in a given month, since it might be equivocal. My baseline assumption is that the expected annual earnings in a given month that people report is an extrapolation from their expected monthly earnings in this given month over the rest of the year (e.g. they multiply this monthly earnings by twelve adjusting for month-specific effects). However, the method is robust to people reporting instead their expected annual earnings in the calendar year of this given month, and to other interpretations of the question. The method is also relatively robust to respondents overestimating the persistence of their earnings as identified in Rozsypal and Schlafmann 2017. Finally, I show in Appendix B.4 that, with this method, a correlation between the persistence ρ and the log-persistent component p_t^i of the respondents, as in Arellano, Blundell, and Bonhomme 2017, would only change the interpretation of the measure.

Ruling out anticipations One possible issue is that reported expected future earnings could include the future transitory component, if the latter is anticipated four months ahead. In that case, what I measure would not be $persist_t^i$ but $persist_t^i e^{\varepsilon_{t+1}/\rho}$, which could bias the results. I show in Appendix B.5 that I can rule it out, because the covariance between my measure of persistent earnings and the realized innovation to log-income at $t + 1$ is small and not significant—while it would be positive and significant if my measure of persistent earnings incorporated some of the innovation that would realize at the next period.

Ruling out independent additive earnings shocks. Another issue would be the presence of additive shocks χ , independent of other shocks, so that: $(y_t)^{true} = y_t + \chi_t$. Note that such shocks would not be a problem if they are i.i.d., but they could generate a bias in the way I build persistent earnings when I take the log of $E_t^i[y_{t+1}^i]$ if their expected value in is non-zero. To look into this, I compute each respondent’s individual-specific variance of future earnings, conditional on future employment. In the absence of independent additive shocks, this variance should scale in the square of persistent earnings (exactly so when $\rho = 1$ and approximately so when ρ is close to one). In the presence of independent additive shocks, an intercept should arise. I show Appendix B.6 that when I regress this variance, detrended from the effect of demographics, over persistent earnings and the square of persistent earnings allowing for an intercept, neither the intercept nor the effect of the level of persistent earnings are significant, while the effect of persistent earnings squared is significant at the 5% level. Thus, the individual-specific variance of future income appears proportional to the square of persistent earnings.

Comparisons with existing results on earnings risk along the earnings distribution. How do these results compare with those of the literature on earnings risk along the earnings distribution? What I find is that, in my sample, the detrended variance of future earnings conditional on future employment is proportional to persistent earnings. This means, among employed people, the variance of the shocks to log-earnings—excluding nonemployment shocks—should be independent of earnings. This is consistent with the stylized fact established by Guvenen, Karahan, Ozkan, and Song 2021 that, among employed people, the standard deviation of log-income growth—including nonemployment shocks—is decreasing in earnings. Indeed, in the earnings specification proposed by Guvenen, Karahan, Ozkan, and Song 2021 to match this stylized fact, the way in which increased earnings reduce the risk to log-income is exclusively through decreased nonemployment risk.

These results are also consistent with the work of Arellano, Bonhomme, De Vera, Hospido, and Wei 2021. This study finds that, among Spanish individuals, the coefficient of variation of future income correlates negatively with current income, with the effect concentrated among the very low income earners and the evolution of future income risk with current income becoming flat quickly. I confirm with an exercise described here (<https://>

[//www.dropbox.com/s/86ml814pxad4fqg/ArellanoRep.pdf?dl=0](https://www.dropbox.com/s/86ml814pxad4fqg/ArellanoRep.pdf?dl=0)) that these two results hold in my dataset. Note that the coefficient of variation of income does not capture unemployment risk, which cancels out if the probability of nonemployment depends mostly on current variables and if income is close to zero when non-employed, so nonemployment risk cannot explain the negative relation. However, consistent with Arellano, Bonhomme, De Vera, Hospido, and Wei 2021’s finding that the result is driven by non-employed and people with low attachment to the labor force, I find that the negative relation disappears when I exclude currently non-employed people and people with very low earnings, for whom the coefficient of variation is very high.

My findings are also consistent with the estimates of Braxton, Herkenhoff, Rothbaum, and Schmidt 2021. The paper additionally finds that persistent earnings risk has risen more among high-skill workers since the 1980s in the US.

4 Empirical effect of persistent earnings on the MPCs

4.1 Specification and results

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

$$MPC_t^i = a_1 + a_2 \text{perst}_t^i(1 + b_2 \text{hh size}_t^i) + a_3 \text{earn}_t^i(1 + b_3 \text{hh size}_t^i) + a_4 \text{wealth cat}_t^i(1 + b_4 \text{hh size}_t^i) + a_5 \text{hh size}_t^i + a_6 \text{state}_t^i + a_7 \text{fixed dem}^i + a_8 \text{period}_t + a_9 \text{age group}_t^i + \xi_t^i, \quad (4.1)$$

with MPC_t^i denoting the reported MPC out of hypothetical shocks of respondent i at period t , perst_t^i the persistent earnings level of this respondent (rescaled so a one dollar change in the persistent component of earnings correspond to a one dollar change in current annual earnings at the average sample values of the transitory component, fixed effect, and time-trend of earnings), hh size_t^i a vector of dummies for the number of members in the household of this respondent, earn_t^i his or her current (total) earnings, wealth cat_t^i a vector of dummies for the wealth category of the household, state_t^i a vector of dummies for its state of residence, and fixed dem^i a vector of dummies for time-invariant demographics (educational attainment and willingness to take risk), period_t a vector of dummies for the four periods over which I have observations (November 2015–November 2018), age group_t^i a vector of age group dummies, and ξ_t^i a noise term.

Importantly, I cannot easily merge the two steps that are, first, building persistent earnings, and, second, estimating their impact on the MPC, into one single step—where I would for instance measure the effect of expected future earnings conditional on employment over the MPC. This is, among other things, because perst_t^i is an exponential of p_t^i . The effect of demographic variables would thus enter twice the specification, linearly and exponentiated and

multiplying the effect of expected future earnings (with extra interactions through $\rho < 1$).

Implementation. I estimate the specification described by (4.1) with a linear regression. The variable $pers_t^i$ is built as described in the previous subsection.

	MPC neg.	MPC pos.
Persistent earnings in \$10,000 (a_2)	0.015*** (0.005)	0.014** (0.006)
Earnings in \$10,000 (a_3)	-0.009** (0.005)	-0.009 (0.006)
One std. dev. persistent earnings	0.048	0.046
Average MPC	0.797	0.545
R^2	0.166	0.236
Observations	1,097	1,113

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

Table 2: Effect of persistent earnings on the MPC

Effect of persistent earnings on the MPC. Table 2 presents selected results from the estimation of specification (4.1). In the first column, the dependent variable is the MPC out of a negative transitory income shock. The first line shows that, conditional on total earnings, wealth, and demographics, an increase in the persistent component of earnings raises this MPC. The point estimate of the effect is 0.015, significant at the 1% level. This means that, comparing individuals with the same total annual earnings, wealth and demographics, people with a \$10,000 higher level of persistent earnings report that their MPC out of a transitory income loss is 0.015 higher (they would cut their consumption by an extra 1.5% of the loss). This is in line with the theoretical prediction of the previous section: everything else being equal, at a higher level of persistent earnings, people respond more to shocks.

The second line shows that, conditional on persistent earnings, wealth, and demographics, an increase in total earnings reduces the MPC out of negative transitory shocks. The point estimate of the effect is -0.009 , significant at the 5% level. This means that, comparing individuals with the same persistent component of earnings, wealth and demographics, people with a \$10,000 higher level of total earnings (which would therefore come from having a higher non-persistent component of earnings), report that their MPC out of a transitory income loss would be lower by 0.009. This is also consistent with what the standard life-cycle model predicts: an increase in the non-persistent component of earnings is akin to an increase in cash-in-hand, which reduces people's sensitivity to shocks.

The next two lines put the magnitude of the main result into perspective. They show that the point estimate of 0.015 implies that a one-standard deviation increase in persistent earnings raises the MPC out of a negative shock by 0.048. Such an increase represents 6% of the average

value of the MPC in the sample, which is 0.796.

The second column shows that these result still hold true when considering the MPC out of a positive transitory income shock. Everything else being equal, persistent earnings affect positively the MPC: the point estimate of the effect is 0.014, significant at the 5% level. The point estimate of the effect of total earnings is also negative, but no longer significant. To put the magnitude into perspective, these results implies that a one standard deviation increase in persistent earnings raises the MPC out of a positive shock by 0.046. This is 8% of the average value of this MPC, which is 0.545.

Since the prediction of the theoretical model is about people with positive net wealth, I re-run the estimation excluding the small amount of people with strictly negative levels of total wealth. I find that the results are very similar, with slightly larger (but not significantly different) point estimates. Table 14 in Appendix C.1 details the results.

Bootstrapping. Although it is commonplace in the consumption insurance literature to use residuals as variables, I further examine the impact of the variability introduced by the first step of generating residuals. To do so, I recompute the standard errors with a bootstrap procedure that includes the first step in the bootstrapping loop. The coefficients remain significant. The only two differences are that the effect of persistent earnings on the MPC out of a positive shock and the effect of total earnings on the MPC out of a negative shock are no longer significant at the 5% level but only at the 10% level. The detailed results are in Table 15 Appendix C.2.

Comparison with a specification that does not consider persistent income separately. To evaluate the importance of treating persistent earnings separately, I also estimate a specification in which I only consider total earnings. I find that the effect of total earnings becomes very small and no longer significant. The point estimates imply that a \$10,000 increase in total annual earnings raises the MPC out of a negative shock by 0.001 and the MPC out of positive shock by zero—and none of these effects are statistically different from zero. The detailed results are in Table 16 Appendix C.3. This result can explain why, although I find a positive and significant effect of persistent earnings on the MPCs, existing studies relying on total earnings mostly show no significant effect (see e.g. Parker, Souleles, Johnson, and McClelland 2013, Boutros 2021, Parker, Schild, Erhard, and Johnson 2022).

4.2 Alternative specifications

Controlling for transitory earnings instead of total earnings. The reason why I control for the total current earnings of the respondents when I measure the effect of persistent earnings on the MPC is because the hypothetical MPC question is framed in terms of percentage of income. Therefore, for the MPCs to be comparable, I need to control for income (which I proxy with earnings): else, because people with a higher earnings are asked to think about larger

shocks, and because the MPC out of positive shocks are known to decrease with the magnitude of the shock (see e.g. Fagereng, Holm, and Natvik 2021 and Golosov, Graber, Mogstad, and Novgorodsky 2021 for natural experiments, and Fuster, Kaplan, and Zafar 2020 and Bunn, Le Roux, Reinold, and Surico 2018 for hypothetical MPCs derived from survey questions) my estimates would be biased downward. This control, however, changes the interpretation of the coefficient associated with persistent earnings: the coefficient captures the combined effect of a one dollar increase in persistent earnings and of the compensating decrease in the rest of earnings—transitory earnings gets multiplied by $\frac{e^{p_i}}{1+e^{p_i}} < 1$ — that keeps total earnings unchanged. Thus, controlling for total earnings has two effects: (i) it makes the MPCs out a percentage change in one’s income comparable across individuals; (ii) it adds to the change in persistent earnings (the one I am interested in) another change, which is a change in transitory earnings (to keep total earnings constant). The first effect is the reason why I do this control, while the second effect is undesirable. Unfortunately, I cannot disentangle between the two. However, to get a sense of their joint importance, I also estimate a model in which I control for transitory earnings instead of total earnings. The results, presented in Table 17 Appendix C.4, show that the effect of persistent earnings on the MPC out of negative shocks is still significant at the 1% level and its point estimate is 0.007. The effect of persistent earnings on the MPC out of positive shocks is no longer significant and its point estimate is 0.006. These numbers constitute a lower bound on the true effect of persistent earnings on the MPC, since this specification is likely downward biased, while my baseline results constitute an upper bound.

Non-linear specification. The specification described by (4.1) allows for a non-linear relation between the level of consumption, wealth and persistent earnings, but it does impose a linear relation between the *MPC* (the partial effect of wealth on consumption), wealth and persistent earnings. Since the theoretical model implies the existence of higher order terms in this relation, I check that the results do not change substantially when I allow for such higher order interactions. Table 18 in Appendix C.5 presents the results of a specification that additionally includes the second order effects of persistent earnings and total earnings on the MPCs, as well as interactions between persistent earnings and the wealth category dummies, between total earnings and the wealth category dummies, and between persistent earnings and total earnings. The average partial effect of persistent earnings on the MPCs out of negative and positive shocks remains significant and positive, with a slightly larger effect of persistent earnings on the MPC out of positive shocks than in the baseline specification though the difference is not significant. The R^2 coefficients increase only by 0.03, from 0.166 to 0.193 and from 0.236 to 0.266, in these more general specifications. In numerical simulations of the model, I later verify that my baseline specification with a linear relation between the MPC, wealth and persistent earnings explains most of the fluctuations in the MPC.

Other measures of persistent earnings. I examine how the effect of persistent earnings on

the MPCs changes when the measure of expected future annual earnings that I use to build persistent earnings is derived from questions about the probability of earnings-related events rather than from a direct question. The results, detailed on Table 19 in Appendix C.6, suggests that the effect is robust to the choice of the measure of expected future earnings. With this alternative measure, the point estimates of persistent earnings on the MPCs out of negative and positive shocks are 0.018 and 0.016, significant at the 5% and 10% level.

I also examine what happens when I build a measure of persistent earnings that captures not just e^{p_t} in 3.1, but also part of the individual fixed effect component e^{α^i} : to do so, instead of regressing $\ln\left(\frac{E_t^i[y_{t+1}^i]}{(1-p_{v_t^i})}\right)$ over education, willingness to take risk, age, gender, and period dummies, I simply regress it over age, gender, and period dummies, to keep the effect of education and willingness to take risk dummies in the residual, thus in the measure of persistent earnings. The results, detailed on Table 20 in Appendix C.6, show that the effect of persistent earnings remains significant at the 1% and 5% on the MPCs out of negative and positive shocks. The point estimates are larger but not statistically different.

Using consumption data rather than hypothetical MPCs to examine the effect of earnings on MPCs. Because my measure of the MPC is based on a hypothetical question about what people would do, rather than what they have done, it might be subject to some biases. I now consider a different specification that uses reported consumption instead of these hypothetical questions to examine the effect of earnings on people’s response to changes in wealth. In this specification, detailed in Appendix C.8, I measure the interaction between the effects of non-housing wealth and persistent earnings on consumption, which is a proxy for the effect of persistent earnings on the MPC: the effect of non-housing wealth on consumption measures a form of MPC, so the interaction measures the influence of persistent earnings on this MPC. However, what it captures still differs from the effect of persistent earnings on the MPC that I measure in the baseline specification for at least three reasons. First, the effect of changes in wealth on consumption are not exactly MPCs out of shocks: wealth changes are not necessarily exogenous and might reflect a response to other events also affecting consumption directly—that is why people rely on natural experiments rather than on regressions of consumption over wealth to measure MPCs. Second, the consumption level is indirectly recovered from other variables thus obtained for only a fraction of the sample, and covers only typical consumption excluding large infrequent purchases. Third, the variations in non-housing wealth are coming from variations of a categorical variables, thus less precise than if the variable had initially been continuous. Despite these limitations, the results, presented in Table 16 in Appendix C.8, show that the interaction between non-housing wealth and persistent earnings is positive and significant, consistent with the results that I obtain using hypothetical MPCs. The magnitude of the effect is larger, as a \$10,000 increase in persistent earnings raises the effect of a one unit increase in non-housing wealth on consumption by 0.027, significant at the 5% level. Also, the

interaction between the effects of non-housing wealth and earnings on consumption is negative and significant.

4.3 Three implications

Bias when measuring the effect of wealth on the MPC without controlling for earnings and persistent earnings. An important implication of my finding is that, persistent earnings and wealth affect the MPC in opposite directions. Because they are also positively correlated—people with higher persistent earnings tend to accumulate more wealth—, this means that measuring the effect of wealth on people’s MPCs without controlling for persistent earnings generates a downward bias: the positive effect of persistent earnings partly confounds the negative effect of wealth on the MPC. This might explain why the effect of wealth on the MPC, which is typically measured without controlling for persistent earnings since the latter is not directly observed, is only significant in recent studies that rely on large datasets capable to detect even small impacts and remains quite modest. 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.’. More precisely, Johnson, Parker, and Souleles 2006 and Parker, Souleles, Johnson, and McClelland 2013 find no significant effect, while the studies of Fagereng, Holm, and Natvik 2021, Baker 2018, or Aydin 2019 find a significant but relatively small effect. In an even more recent work, Ganong, Jones, Noel, Farrell, Greig, and Wheat 2020 find that moving from the lowest level of liquid asset to the highest one reduces the MPC by 0.27 point. This is also modest given the magnitude of the change in liquid assets required for a 0.27 change in the MPC.

In the SCE, since I can build a measure of persistent earnings, I can estimate the effect of liquid wealth on the MPC both without and with controls for persistent earnings and total earnings. Table 3 presents the results. The first and second columns report the coefficients associated with the wealth category dummies, when estimating a version of equation (4.1) without the persistent earnings and earnings variables. The reference category for wealth is ‘Less than \$500 assets’, so the coefficient associated with a given category capture the effect on the MPC of shifting from less than \$500 in non-housing wealth to the category considered. These first two columns show that, absent any controls, having more wealth than \$500 generally associates with a higher MPC out of a negative transitory shock and does not correlate with a different MPC out of a positive shock, as the literature based on small survey data finds.

The third and fourth columns report the coefficients associated with the wealth category dummies, when estimating equation (4.1), thus including the effect of persistent earnings and total earnings. In contrast to the first column, the coefficients are no longer positive and significant for the MPC out of negative shocks, but generally non-significant or negative and significant, in particular when considering a shift to high levels of wealth. In contrast to the second

	Without any earnings control		With earnings control	
	MPC neg.	MPC pos.	MPC neg.	MPC pos.
Less than \$500 assets
\$500-\$999 assets	0.139*** (0.139)	0.186 (0.186)	0.047 (0.047)	0.06 (0.06)
\$1,000-\$1,999 assets	0.863*** (0.129)	0.664*** (0.185)	-0.064 (0.055)	-0.114* (0.062)
\$2,000-\$4,999 assets	0.896*** (0.126)	0.256 (0.42)	0.015 (0.042)	-0.089 (0.054)
\$5,000-\$9,999 assets	0.144 (0.165)	-0.309* (0.163)	0.009 (0.043)	-0.127** (0.055)
\$10,000-\$19,999 assets	1.009*** (0.121)	-0.29 (0.177)	-0.01 (0.043)	-0.246*** (0.054)
\$20,000-\$29,999 assets	0.645*** (0.154)	0.065 (0.2)	-0.033 (0.05)	-0.196*** (0.059)
\$30,000-\$49,999 assets	0.832*** (0.123)	0.698*** (0.173)	-0.055 (0.047)	-0.258*** (0.058)
\$50,000-\$99,999 assets	0.425*** (0.133)	-0.168 (0.183)	-0.081* (0.047)	-0.319*** (0.056)
\$100,000-\$249,999 assets	0.839*** (0.103)	0.304 (0.232)	-0.044 (0.049)	-0.383*** (0.056)
\$250,000-\$499,999 assets	0.142 (0.107)	0.034 (0.157)	-0.148** (0.058)	-0.245*** (0.075)
\$500,000-\$749,999 assets	0.588*** (0.113)	0.001 (0.138)	-0.056 (0.084)	-0.499*** (0.065)
\$750,000-\$999,999 assets	-0.285 (0.193)	-0.248 (0.175)	-0.443*** (0.086)	-0.676*** (0.086)
More than \$1,000,000 assets	0.071 (0.105)	-0.157 (0.144)	0.002 (0.09)	-0.369*** (0.108)
R^2	0.155	0.226	0.166	0.236
Observations	1,108	1,125	1,097	1,113

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

Table 3: Effect of wealth on the MPC without and with earnings controls

column, the coefficients are no longer non-significant for the MPC out of positive shocks, but generally negative and significant. This is consistent with the existence of a substantial downward bias that arises when estimating the effect of wealth on the MPC without controlling for the effect of persistent earnings and earnings. The coefficients are almost unchanged when controlling only for persistent earnings and not for total earnings (Table 25 Appendix D.1).

Limited gains from a narrow targeting of fiscal stimulus to low income people. Over the past two decades, policy makers increasingly relied on fiscal stimulus in the form of direct payments made to households. In the US, those payments have not been too narrowly targeted towards very low-income people: in 2008 and 2020, the payments that people received

decreased only around the 10th percentile of the adjusted gross income distribution.²⁰ Some argued that a more narrow targeting of the payments could deliver a stronger consumption response. Those payments were also conditional on a minimum income of \$ 3,000 per year, so it makes sense to consider the impact of targeting within my sample of employed people with earnings above \$ 1,885.

My result that persistent earnings raise the MPC suggests that targeting more narrowly low income people may not have large effects on the average MPC out of stimulus payments: although people with lower total income are likely to have a lower level of wealth and thus a higher MPC, they are also likely to have lower persistent earnings, which might partly counteract the effect of wealth on their MPC.

	Average MPC neg.	Average MPC pos.
Earnings < 10th	0.806	0.591
Earnings < 25th	0.813	0.600
Earnings < 50th	0.809	0.591
Earnings < 75th	0.798	0.579
Earnings < 90th	0.793	0.569
All	0.788	0.555
Wealth < 13th	0.849	0.747
Wealth < 13th & Perst. > 50th	0.861	0.807

Table 4: Effect of targeting on average MPC

I examine whether this implication of my results is true in the SCE. Table 4 presents the average MPC among different groups of people. The first six lines show that the average MPCs out of negative and positive shocks are not substantially higher among people with lower earnings. More precisely, comparing those with earnings below the 10th percentile to those with earnings below the 25th percentile, the average MPCs are larger in the second group (from 0.806 to 0.813 out of negative shocks and from 0.591 to 0.600 out of positive shocks). The average MPCs then decrease modestly when targeting broader groups of low-income people. Absent any targeting, the average MPCs are still 0.788 and 0.555, not much smaller than the average value among people below the 10th or the 25th percentile. Given (i) that this average MPC reporting is already biased towards being higher for people with lower income—because those people are asked about their response to a shock of a smaller magnitude while the MPC out of smaller payments is typically higher²¹—, and (ii) that narrower targeting would also mean giving out larger payments while, again, larger payments reduce the MPCs, the actual differences in average consumption responses across targeting strategies are likely to be even

²⁰For both the 2008 and 2020, the payments were phased out for taxpayers with adjusted gross income greater than \$75,000 (\$150,000 for couples filing jointly) in the previous year.

²¹See e.g. Fagereng, Holm, and Natvik 2021 and Golosov, Graber, Mogstad, and Novgorodsky 2021 for natural experiments, and Fuster, Kaplan, and Zafar 2020 and Bunn, Le Roux, Reinold, and Surico 2018 for hypothetical MPCs derived from survey questions.

smaller than what Table 4 shows.

The seventh and eighth lines suggest that the policy makers could target a group of people with substantially higher MPCs than others if they could observe non-housing wealth, or both non-housing wealth and the persistent component of earnings: the average MPC values rise from 0.79 and 0.56 in the whole sample to 0.85 and 0.75 among people with non-housing wealth below \$1,000; they rise to 0.86 and 0.81 among people with non-housing wealth below \$1,000 and persistent earnings above the median. However, both wealth and persistent earnings are difficult to observe for policy makers, thus difficult to use to condition payments. In addition, as discussed in the paragraph above, it is still true that comparing these averages is likely to yield starker differences than the policy would, so even with these more substantial differences it is hard to establish that targeting payments would yield much stronger average consumption responses.

Policy makers might however want to target the fiscal stimulus to certain groups of people for reasons other than raising the average MPC out of the payments. Furthermore, although the result from which I derive this implication concerns employed people—and thus only have direct implications for them—I also look at the effect of targeting low-income people when I include both nonemployed and employed (even those with earnings below \$1,885 per year) in the sample and consider percentiles of income (earnings plus UI benefits) rather than earnings. Appendix D.2 presents the results. The increase in average MPC is a little more pronounced.

Similarities between the average MPCs of homeowners with low and high income. Finally, my main result implies that even homeowners with high income—who, as the new narrative of the Great Recession suggests, bore a substantial part of the wealth loss at the onset of the crisis—can respond strongly to a transitory income loss or to a wealth loss, thus contribute to the plunge in consumption that followed the wealth loss at the beginning of the Great Recession.

	Average MPC neg.	Average MPC pos.
Homeowners with earnings < 50th	0.773	0.532
Homeowners with earnings \geq 50th	0.768	0.480

Table 5: Average MPCs of homeowners with low and high earnings

To examine whether this is true in the SCE, I compute the average MPCs of homeowners below and above the average earnings. I compute homeownership status from people’s answers to questions about their housing wealth.²² Table 5 presents the results. They show that the average MPC out of negative shocks is almost identical in the two groups (equal to 0.77). It confirms that homeowners with high income can be as responsive to a negative wealth shock

²²More precisely I label as homeowners people who report a strictly positive level of housing wealth and also give a response to the question about liquid wealth (just to make sure they answer this module of the questionnaire seriously).

as homeowners with low income (likely to be subprime borrowers). The average MPC out of positive shocks is still a little higher among homeowners with earnings below median than among those with earnings above median.

5 Comparison with simulated data

5.1 Model and calibration

Consumers' maximization problem. I simulate and calibrate a rich life-cycle model that mimics the situation of US households, to understand whether such a life-cycle model is consistent with the survey data estimates, and to examine what channels are quantitatively the most important for the effect of persistent earnings on the MPC in the simulated model. A period is a year, since this is the timespan that people are asked to consider in the survey. The period utility $u(\cdot)$ is a log-utility function. I add a minimum consumption threshold, so it is not possible to consume less than \$2,175 per year and marginal utility approaches infinity as consumption approaches this threshold. People maximize their utility weighted by a function of demographic characteristics $e^{\delta_t z_t}$ at each period t , with z_t a vector of demographics. There are no changes in demographics until age 49 ($e^{\Delta \delta_t z_t} = 1$ for everyone). After that, each year, the change in demographics is such that $e^{\Delta \delta_t z_t} = 0.985$. This is to match the hump-shaped pattern of consumption over the life-cycle, which Attanasio, Banks, Meghir, and Weber 1999 and Attanasio 1999 document and show to be disappearing when controlling for demographics. Intuitively it should capture that people are done paying for some expenses that are life-cycle specific (e.g. children-related including college or work-related with a documented shift around retirement in Aguiar and Hurst 2005, Aguiar and Hurst 2007, and Hurd and Rohwedder 2013). I choose age 49 because that is the shifting point in the hump-shaped patterns documented in Attanasio, Banks, Meghir, and Weber 1999 and Attanasio 1999.²³

Wealth. I assume that wealth in the model represents net risk-free liquid wealth. This means assuming that people may have illiquid wealth on the side, but that they do not use it to smooth consumption.²⁴ Kaplan and Violante 2014 establish that illiquid wealth (e.g. housing) is rarely liquidated to smooth out consumption. Kaplan and Violante 2022 show that a one-asset model that matches the level of liquid wealth that people hold, rather than their total wealth, can generate MPCs consistent with empirical evidence. Thus, while another, more realistic way to generate MPCs that match the empirical evidence is to have a model with two assets, one liquid

²³See Figure 1 in Attanasio, Banks, Meghir, and Weber 1999 and Figure 4 in Attanasio 1999.

²⁴The model is for instance equivalent to one in which people would start their working life with a house and an exogenous amount of mortgage they have to repay by fixed amounts at each period. Their mortgage payment would be the equivalent of a rent, that is, the price of a housing service—mortgage payments are expenses that I take into account in the data, as rent is. When people die, they pass on a fraction of the house to their children, who take a mortgage to buy it full and renovate it.

and one illiquid, the one-asset model works well when the objective is to model consumption.

The yearly interest rate on the liquid asset is constant and set to $r = 0.01$, to match with the low interest rate on liquid holdings over the period 2015-2018.²⁵

To measure the average level of net liquid wealth that people have in the survey data, I use the detailed questions from the Household Finance survey, with a correction to account for the fact that the questions in the survey are about the wealth and debt of the respondents and their spouses while the households are all single in the model. Net liquid wealth corresponds to liquid wealth (checking and savings accounts, stocks, bonds, mutual funds, Treasury bonds) minus non-housing debt.²⁶ The average amount of liquid wealth I obtain in the data is \$ 3,561.

Discount factor. I set the discount factor to match the average level of net liquid wealth in the survey data (+/-1%). The value that matches it in the baseline framework is $\beta = 0.951$.

Borrowing limit. In addition to the period budget constraints, people face a borrowing limit. In the baseline calibration, I fix it at a maximum debt of \$3,261—roughly consistent with the magnitudes of the 2021 SCF in which 45% of US households report revolving balances on one or more of their credit cards at the time of the survey, and the median revolving family owes \$2,700.

Lifespan. People enter the labor market at age 25. They retire at age 62. After retirement, they have a non-zero probability to die at each period from age 62 to age 91. The survival probabilities are obtained from the National Center for Health Statistics (I use the data from Kaplan and Violante 2010). If still alive at age 91, a household dies with certainty at age 92.

Earnings. The earnings that people get at each period follows exactly the parametric process proposed in Guvenen, Karahan, Ozkan, and Song 2021. As such this process is a particular case of the specification (3.1)-(3.3) I assume in the empirical section—except for making the probability of unemployment depend on the contemporaneous value of persistent earnings rather

²⁵Incidentally, since the discount factor β is set to match the empirical level of liquid wealth, changing the interest rate leads to an adjustment in β and has little impact on the simulation results.

²⁶This is a little more liquid than the measure I use in the estimation of the main specification because it excludes retirement wealth (defined contribution plans and individual retirement accounts). The reason for the difference is that, using precise questions from the SCE, I have the possibility to compute truly liquid wealth, which is my preferred measure of wealth, although I do not observe it for enough respondents in my sample to use it empirically. To control for the fact that in the survey people are not single, I multiply the observed level of liquid wealth minus non-housing debt by 0.7480, which I obtain as a weighted average of a share (2/3) for households with at least two adults and 1 for households with only one adult. I compute this over the 590 people in my main sample (with at least one the two MPCs, categorical wealth, persistent earnings, and time-varying demographics observed) for whom I also observe these variables.

than on its previous value:

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

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

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

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

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

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

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

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

This means that, compared to a simple transitory-persistent model with two state variables—the current persistent and the current transitory components—, there are two additional variables I need to observe to determine earnings: the realization of the fixed-effect and the employment status. I calibrate the parameters of this process from the estimates of Guvenen, Karahan, Ozkan, and Song 2021 (summarized in Appendix E.1 of this paper and taken from Table IV of their paper and Table D.III of their Online Appendix).

Transfers, taxes, and social security income. People also receive transfers that keep their annual income at a minimum of \$2,175: if their earnings realization falls below this threshold (either because they are not employed or because their earnings are lower than this threshold), the government gives them the difference required to keep their revenue at this threshold. The choice of this threshold is partly motivated by the fact that Guvenen, Karahan, Ozkan, and Song 2021 select out people with earnings below \$1885 per year (in \$2010)—which corresponds one quarter of full-time work at half of the minimum wage in 2010—and Gouveia and Strauss 1994 exclude people with earnings below \$3000 per year (in current dollars over the period 1979-1989). These transfers also insure that people are always able to meet the minimum consumption threshold of \$2,175.

People then pay taxes according to the nonlinear tax function of Gouveia and Strauss 1994, $\text{tax}(y_t^i) = \tau^b (y_t^i - ((y_t^i)^{-\tau^p} + \tau^s)^{-1/\tau^p})$ parametrized with $\tau^b = 258$, $\tau^p = 0.768$, $\tau^s = 2.0e - 4$ as in Kaplan and Violante 2010.²⁷

²⁷Contrary 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 is pre-tax earnings and I use the tax function to recover net earnings. People pay

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

Transitory income shock. To compute people's MPCs, I simulate a positive and a negative wealth shock that corresponds to 10% of their current income, as in the survey question. The shock occurs once for each individual and the timing of the shock is uniformly distributed between age 26 and age 55.

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 2006.²⁸

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. For comparability with my survey data, I convert my simulated data from 2010\$ to 2014\$.²⁹

Selection. I select individuals age 25-55 and employed at the moment when they experience the transitory shock. I trim the top and bottom 1% of the persistent component of earnings. In comparison, in the survey data, I trim expected future earnings (and other variables that I do not use in this section) because I do not directly observe the persistent component.

Model fit. Table 6 presents the average level of net liquid wealth, annual earnings, and MPCs among employed individuals in the survey data and in the simulations. In the survey data, these are computed over the individuals that are in the final samples over which I run my main regression (either in the MPC negative or MPC positive sample); in the simulated data, these are computed over the individuals that remain after selection.

taxes on their income, which here is the sum of their earnings and transfers, but the tax rate is zero (or quasi zero) at the income level that opens rights to transfers.

²⁸The number of grid points is as follows: the grid for wealth has 150 exponentially spaced grid points; the grid for the persistent component of earnings is age-varying and at each age has 35 equally spaced points; the grid for the transitory shock has 11 equally spaced points; the grid for the fixed effect component of earnings has 9 equally spaced points; the grid for lifetime average earnings (used to compute retirement income) has 9 equally spaced points. Expanding the grid further does not change the results.

²⁹This is why the value of the expenditures threshold, transfer and borrowing limit are not round numbers: in 2010\$ the threshold and transfer are \$2,000 and the borrowing limit is \$3,000.

	Survey data	Simulated data
Net liq. wealth	3,561	3,531
Earnings	63,471	58,251
MPC neg.	0.797	0.663
MPC pos.	0.545	0.579

Table 6: Model fit

The first line illustrates my targeting the average amount of net liquid wealth in the data to calibrate the discount factor β in the model. As a result, the amount of wealth in the simulations with a liquid wealth calibration matches at $\pm 1\%$ the liquid wealth in survey data.

The second line indicates that the average earnings in the survey data is \$63,471. This is relatively close to the simulated data average of \$58,251. This latter amount is exclusively generated by the earnings process proposed in Guvenen, Karahan, Ozkan, and Song 2021, calibrated from the parameters they estimate in administrative data, and converted in \$2014. It means that the earnings people report in survey data resembles the earnings generated from a process designed to fit administrative data.

The third and fourth line presents the average MPCs out of negative and positive shocks. In the survey data, both MPCs are large, with an asymmetry between the two: 0.797 for a negative shock and 0.546 for a positive shock. In the simulated data, the MPCs are large as well, although large MPCs are notoriously difficult to generate in life-cycle models. There is also a marked asymmetry between the two, but a little smaller than in the survey data: the average MPC out of a negative shock is 0.663 and the average MPC out of a positive shock is 0.579.

The sources of the high average MPCs. In Appendix E.2, I show that two necessary ingredients to generate high MPCs out of both negative and positive shocks are incorporating a rich earnings process (with individual fixed effects, unemployment shocks, transitory and persistent shocks drawn from a mixture of normal distributions, and a dependency between persistent earnings and the probability of non-employment) rather than a simple transitory-persistent process, and doing a liquid wealth calibration (the assumption that people cannot easily draw on their illiquid wealth to finance their consumption) rather than a total wealth calibration. Other features of the model, the demographic trend after age 49, the consumption threshold and transfers, or the strict borrowing limit can be removed without substantially modifying the MPCs.

5.3 Comparison between survey and simulated data estimates

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

employed people.

Specification. I build a set of wealth category dummies with the same thresholds as in the survey question. The estimation equation is similar to (4.1), with again the difference that the only demographic variables are the year (equivalently age) dummy variables. I create one dummy for each year (equivalently each age), instead of having broader period (or age) groups. The equation that I estimate is then:

$$MPC_t^i = a_1 + a_2 \text{perst}_t^i + a_3 \text{earn}_t^i + a_4 \text{wealth cat}_t^i + a_7 t + \xi_t^i. \quad (5.9)$$

	Survey data (SCE)		Simulated data	
	MPC neg.	MPC pos.	MPC neg.	MPC pos.
Pers. earn. in \$10,000	0.015	0.014	0.014	0.010
Earn. in \$10,000	-0.009	-0.009	-0.002	-0.003
One std. dev. persistent earnings	0.048	0.046	0.055	0.040
Av. MPC	0.797	0.545	0.657	0.579
R^2	0.166	0.236	0.797	0.754
Observations	1,097	1,113	3,478	3,469

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

Table 7: Effect of persistent earnings on the MPC in survey data and in simulations

Results. Table 7 presents the effect of persistent earnings and earnings on the MPC. The first two columns are a remainder of the results I obtain in the SCE. The third and fourth columns present the results I obtain in data simulated as described above.

The two sets of results are similar. The third column shows that, in the simulated data, a \$10,000 increase in persistent earnings raises the MPC out of a negative transitory income shock by 0.014, which is similar to the increase of 0.015 that I estimate in survey data. Conditional on persistent earnings, total earnings has a negative impact on the MPC, although the point estimate is smaller (-0.002 instead of -0.009). These estimates also imply that a one-standard deviation increase in persistent earnings is associated with a 0.055 increase in the MPC. This is close to the 0.048 increase that I estimate in survey data.

The fourth column shows that, in the simulated data, a \$10,000 increase in persistent earnings raises the MPC out of a positive transitory income shock by 0.010. It is a little smaller but still close to the coefficient of 0.014 that I estimate in survey data. The fact that the effect of persistent earnings on the MPC is a little smaller when considering a positive shock than when considering a negative shock is as in the survey results. Conditional on persistent earnings, the effect of earnings is negative, as in the survey results. These estimates imply that a one standard deviation increase in persistent earnings associates with a 0.040 increase in the MPC out of a positive shock, which is close to the 0.046 increase implied by my survey data estimates.

Incidentally, the R^2 of the estimation are large in the simulated data, which confirms that the simple linear relation between the MPC and persistent earnings that I estimate captures most of the fluctuations in the MPC in this standard model.

Disentangling the different channels. Since the model I simulate is richer than the standard model in which I prove that persistent earnings raises the MPC, it is important to examine the role of these additional elements in the results. To do so, I simulate a simpler version of the model that eliminates any element not present in the standard model. More precisely, in this simpler version, I set the persistence ρ to one, I eliminate the dependency between current persistent earnings and the probability of non-employment, re-adjusting with a constant so the average annual earnings is the same as in the simulations of the baseline model, I remove the borrowing constraint, I set the social security income that people receive during retirement to a fixed value that is the same for everyone and corresponds to the average social security income in the simulations of the baseline model, and I set the taxes and minimum transfers at zero—I also set the consumption threshold at zero else some people might not be able to consume above the threshold. I then add each of the extra elements back into the simple standard model. In each case, the discount factor β is re-calibrated so the average wealth is at \$3,561 (+/-1%).

	Baseline model		Simple standard model	
	MPC neg.	MPC pos.	MPC neg.	MPC pos.
Pers. earn. in \$10,000	0.014	0.010	0.021	0.018
Earn. in \$10,000	-0.002	-0.003	0.003	0.000
	Std + persistence below one		Std + unemp dep. on PE	
	MPC neg.	MPC pos.	MPC neg.	MPC pos.
Pers. earn. in \$10,000	0.023	0.017	0.038	0.034
Earn. in \$10,000	0.002	0.002	-0.004	-0.007
	Std + borrowing limit		Std + taxes, transfers, & SSI	
	MPC neg.	MPC pos.	MPC neg.	MPC pos.
Pers. earn. in \$10,000	0.010	0.021	0.005	0.012
Earn. in \$10,000	-0.002	-0.001	0.007	0.000

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

Table 8: Effect of persistent earnings on the MPC

Table 8 displays the results. The first two columns of the first panel present the results of the full baseline model, for comparison. The third and fourth columns of this first panel present the results I obtain when moving to a simpler version of this model, that the standard model described in the theoretical section encompasses. As predicted by theory, these columns show that persistent earnings still associate positively with the MPC. The point estimates are one third to a half larger than in the baseline model: they imply that a \$10,000 increase in persistent earnings raises the MPC out of negative and positive shocks by 0.021 and 0.018. The

effect of the other components of earnings is small and non-significant. Although I recalibrate the discount factor β and the earnings process to keep the average level of wealth and annual earnings the same as in the baseline, the distributions around those averages might be different in the simulations of the simple standard model and in the simulations of the baseline model. As a result, the change in the average effect of persistent earnings and earnings can be due to the effect of persistent earnings conditional on wealth and persistent earnings being larger in the simple standard model, but also to the distributions of wealth and earnings around their average having changed in a way that raises the average effect.

The first and second columns of the second panel present the results I obtain when adding to the simple standard model a persistence $\rho = 0.991$ below one—and re-adjusting so the average annual earnings matches the baseline. This has only a small impact on the results, although the asymmetry between the effect of persistent earnings on the MPCs out of negative and positive shocks increases slightly. The third and fourth columns of the second panel present the results I obtain when adding to the simple standard model the positive relation between persistent earnings and the probability to be employed—and re-adjusting so the average annual earnings matches the baseline. This raises strongly the average effect of persistent earnings on the MPC: the point estimates of the effect of persistent earnings on the MPCs out of negative and positive shocks move from 0.021 and 0.018 to 0.038 and 0.034. Once again, this might be both because of a change in the effect of persistent earnings conditional on wealth and persistent earnings, or because of a change in the distributions of wealth and persistent earnings.

The third panel show that adding a borrowing limit or adding progressive taxes, transfers, and a retirement social security income based on persistent earnings to the simple model has a similar impact on the effect of persistent earnings on the MPC. Both reduce the effect of persistent earnings on the MPC out of a negative shock quite substantially but not the MPC out of a positive shock (or not as much). A possible interpretation is that those elements introduce a kink in the consumption function below which the slope is fixed and no longer affected by persistent earnings, so negative shocks that are more likely to bring people to this part of the consumption function generate a consumption response that is less influenced by persistent earnings. As before, part of the variation in average effects might be due to variations in the distributions of wealth and persistent earnings in these simulations.

Overall, the direct effect of persistent earnings on the MPC, that is, the effect that I identify in the standard life-cycle model and that plays through persistent earnings strengthening precautionary behavior because it directly multiplies current and uncertainty future earnings, is large in numerical simulations calibrated to mimic US households. Other channels, in particular the dependency of the non-employment probability on persistent earnings and the tax and transfer system, play a quantitatively important role in the way persistent earnings affect the MPCs but they partly offset each other. They overall reduce the effect by one third.

6 Conclusion

In this paper, I establish the unintuitive theoretical result that, in the standard life-cycle model used throughout macroeconomic studies, people with a higher persistent component of earnings have a higher MPC, everything else being equal that is conditional on wealth, other earnings components, and demographic characteristics.

While it is important per se to know about the mechanisms at play in this widely used model, I also examine the empirical validity of this theoretical prediction. I find that the prediction holds true in the New York Fed Survey of Consumer Expectations, in which a one standard deviation increase in persistent earnings associates with a 0.05 increase in the reported MPC out of an hypothetical transitory income shock.

Finally, I show that this empirical evidence is also quantitatively consistent with a rich life-cycle model calibrated to mimic the survey data: in numerical simulations of such a model, the MPC levels are close to the ones I observe in survey data, and the effect of persistent earnings on the MPCs similar to the one I measure in survey data. The direct effect I identify theoretically is large in these simulations, and its magnitude is reduced by the additional channels at play in the rich version of the model.

References

- Adelino, M., Schoar, A., and Severino, F. (2016).** “Loan Originations and Defaults in the Mortgage Crisis: The Role of the Middle Class.” *Review of Financial Studies*, 1635–1670.
- Aguiar, M. and Hurst, E. (2005).** “Consumption versus Expenditure.” *Journal of Political Economy* 113.5, pp. 919–948.
- Aguiar, M. and Hurst, E. (2007).** “Life-Cycle Prices and Production.” *American Economic Review* 97.5, pp. 1533–1559.
- Albanesi, S., De Giorgi, G., and Nosal, J. (2017).** *Credit Growth and the Financial Crisis: A New Narrative*. Working Paper 23740. National Bureau of Economic Research.
- Arellano, M., Bonhomme, S., De Vera, M., Hospido, L., and Wei, S (Sept. 2021).** *Income Risk Inequality: Evidence from Spanish Administrative Records*. Working Papers 2136. Banco de España.
- Arellano, M., Blundell, R., and Bonhomme, S. (2017).** “Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework.” *Econometrica* 85, pp. 693–734.
- Armantier, O, Topa, G, Klaauw, W Van der, and Zafar, B. (2017).** “An overview of the Survey of Consumer Expectations.” *Economic Policy Review* 23-2, pp. 51–72.
- Attanasio, O. (1999).** “Consumption.” *Handbook of Macroeconomics*. Ed. by Taylor, J. B. and Woodford, M. Vol. 1. Handbook of Macroeconomics. Elsevier. Chap. 11, pp. 741–812.
- Attanasio, O., Banks, J., Meghir, C., and Weber, G. (1999).** “Humps and Bumps in Lifetime Consumption.” *Journal of Business and Economic Statistics* 17.1, pp. 22–35.
- Auclert, A. (2019).** “Monetary Policy and the Redistribution Channel.” *American Economic Review* 109.6, pp. 2333–67.
- 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, A., Duernecker, G., Forstner, S., and Goensch, J. (2021).** *The Effects of Biased Labor Market Expectations on Consumption, Wealth Inequality, and Welfare*. CESifo Working Paper Series 9326. CESifo.
- Blundell, R., Pistaferri, L., and Preston, I. (2008).** “Consumption Inequality and Partial Insurance.” *American Economic Review* 98.5, pp. 1887–1921.
- Boutros, M. (2021).** “Household Finances and Fiscal Stimulus in 2008.” *mimeo*.
- Braxton, J. C., Herkenhoff, K., Rothbaum, J., and Schmidt, L. (2021).** *Changing Income Risk across the US Skill Distribution: Evidence from a Generalized Kalman Filter*. Working Paper 29567. National Bureau of Economic Research.
- Broda, C. and Parker, J. (2014).** “The Economic Stimulus Payments of 2008 and the aggregate demand for consumption.” *Journal of Monetary Economics* 68.S, pp. 20–36.
- Bunn, P., Le Roux, J., Reinold, K., and Surico, P. (2018).** “The consumption response to positive and negative income shocks.” *Journal of Monetary Economics* 96, pp. 1–15. ISSN: 0304-3932.
- Carroll, C. (1997).** “Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis.” *The Quarterly Journal of Economics* 112.1, pp. 1–55.
- Carroll, C. (2006).** “The method of endogenous gridpoints for solving dynamic stochastic optimization problems.” *Economics Letters* 91.3, pp. 312–320.
- 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. (2011).** *Theoretical foundations of buffer stock saving*. CFS Working Paper Series 2011/15. Center for Financial Studies (CFS).
- Carroll, C. and Kimball, M. (1996).** “On the Concavity of the Consumption Function.” *Econometrica* 64.4, pp. 981–92.
- Carroll, C. and Kimball, M. (2006).** “Precautionary Saving and Precautionary Wealth.” Entry for *The New Palgrave Dictionary of Economics*, 2nd Ed.
- Drèze, J. and Modigliani, F. (1972).** “Consumption decisions under uncertainty.” *Journal of Economic Theory* 5.3, pp. 308–335.
- Fagereng, A., Holm, M., and Natvik, G. (2021).** “MPC Heterogeneity and Household Balance Sheets.” *American Economic Journal: Macroeconomics* 13.4, pp. 1–54.
- Foote, C., Loewenstein, L., and Willen, P. (Aug. 2020).** “Cross-Sectional Patterns of Mortgage Debt during the Housing Boom: Evidence and Implications.” *The Review of Economic Studies* 88.1, pp. 229–259. ISSN: 0034-6527.
- 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, M., Graber, M., Mogstad, M., and Novgorodsky, D. (2021).** *How Americans Respond to Idiosyncratic and Exogenous Changes in Household Wealth and Unearned Income*. Working Paper 29000. National Bureau of Economic Research.
- Gong, L., Zhong, R., and Zou, H. (2012).** “On the concavity of the consumption function with the time varying discount rate.” *Economics Letters* 117.1, pp. 99–101.
- Gouveia, M. and Strauss, R. (1994).** “Effective Federal Individual Tax Functions: An Exploratory Empirical Analysis.” *National Tax Journal* 47.2, pp. 317–39.
- Guvenen, F., Karahan, F., Ozkan, S., and Song, J. (2021).** “What Do Data on Millions of U.S. Workers Reveal About Lifecycle Earnings Dynamics?” *Econometrica* 89.5, pp. 2303–2339.
- Hurd, M. and Rohwedder, S. (2013).** “Heterogeneity in spending change at retirement.” *The Journal of the Economics of Ageing* 1, pp. 60–71.
- Jappelli, T. and Pistaferri, L. (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., Mitman, K., and Violante, G. (2020).** “The Housing Boom and Bust: Model Meets Evidence.” *Journal of Political Economy* 128.9, pp. 3285–3345.
- 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. and Violante, G. (2022).** *The Marginal Propensity to Consume in Heterogeneous Agent Models*. Working Paper 30013. National Bureau of Economic Research.
- Karger, E. and Rajan, A. (May 2020).** *Heterogeneity in the Marginal Propensity to Consume: Evidence from Covid-19 Stimulus Payments*. Working Paper Series WP 2020-15. Federal Reserve Bank of Chicago.
- Kimball, M. (1990a).** “Precautionary Saving and the Marginal Propensity to Consume.” mimeo.
- Kimball, M. (1990b).** “Precautionary Saving in the Small and in the Large.” *Econometrica* 58.1, pp. 53–73.
- Kueng, L. (2018).** “Excess Sensitivity of High-Income Consumers.” *The Quarterly Journal of Economics* 133.4, pp. 1693–1751. ISSN: 0033-5533.
- Lewis, D., Melcangi, D., and Pilossoph, L. (Nov. 2019).** “Latent Heterogeneity in the Marginal Propensity to Consume.” 902.
- 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., Schild, J., Erhard, L., and Johnson, D. (2022).** *Household Spending Responses to the Economic Impact Payments of 2020: Evidence from the Consumer Expenditure Survey*. Working Paper 29648. National Bureau of Economic Research.
- Parker, 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.
- Pistaferri, L. (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. (May 2017).** *Overpersistence Bias in Individual Income Expectations and its Aggregate Implications*. CEPR Discussion Papers 12028. C.E.P.R. Discussion Papers.
- Straub, L. (2019).** “Consumption, Savings, and the Distribution of Permanent Income.” mimeo.
- Toda, A. (2021).** “Necessity of hyperbolic absolute risk aversion for the concavity of consumption functions.” *Journal of Mathematical Economics* 94.C.