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, Persistent Earnings, 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. I am grateful to Mons Chan, Gaetano Gaballo, Jean-Marc Robin, and Fabien Tripier for helpful comments and suggestions.

1 Introduction

How does 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] Researchers have examined how differences in accumulated assets, which I refer to as wealth, and in their liquidity, can produce such heterogeneity, but differences in earnings have not received the same attention. Yet, they might also generate some heterogeneity in consumption responses. Furthermore, the effect of earnings on consumption responses should 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 Auclert 2019 finds that the (unconditional) covariance between income and consumption responses matters for the design of monetary policy. Finally, the effect of earnings on consumption responses can be of interest to understand 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 loss[2], 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, some finding a significant and negative effect, and 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

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

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

[3]Most 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.

components separately. More precisely, earnings are typically modeled with at least two components, 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 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 yearly 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; the model is richer than the one I examine theoretically, but the additional channels of the rich model overall reduce the effect of persistent earnings on the response of consumption.

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 evolve as a random walk. It encompasses as a special case the situation in which persistent earnings is a time-invariant individual component (e.g. a measure of people's inner skills). I show that, for somebody with positive net 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, a change in cash-in-hand leads people to adjust their current consumption for two reasons: first, an increase in cash-in-hand raises people's total expected lifetime resources; second, an increase in cash-in-hand 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, and vice-versa for a decrease in wealth. 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 an increase in persistent earnings strengthens the response of precautionary saving is that it scales up the uncertain part of total resources (future earnings) but not all of certain part (current earnings but not wealth). 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 non-transitory component of earnings be persistent, 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 in these 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]

[4] More 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.

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, 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 my sample 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.

The additional channels through which persistent earnings may effect precautionary behavior and therefore the MPC in this rich model are the fact that persistent earnings modifies the probability of future unemployment, that it is not permanent thus does not multiplies future earnings one-for-one, that it influences the taxes, transfers, and retirement income that people pay and receive, and that 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 mechanism of the standard model plays an important role in the positive relation I obtain between persistent earnings and the MPCs in these simulations.

2 Persistent earnings and the MPC in life-cycle models

2.1 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:

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

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

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

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

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

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

Utility is time-separable and at each period depends only on contemporaneous consumption. The period utility function $u(.)$ is such that marginal utility is positive, decreasing, and convex in consumption: $u'(.) > 0$, $u''(.) < 0$, and $u'''(.) > 0$. Utility also approaches infinity as consumption approaches zero. The discount factor β captures how much a consumer i discounts utility between two consecutive periods. The positive consumption condition (2.2) imposes that consumption be strictly positive at each period. At each period, the household earns the stochastic amount y_t^i. The budget constraint (2.3) states that to store their wealth from one period to another the consumers only have access to a risk-free, liquid asset, with a_t 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_t. To simplify the presentation of the mechanism, I assume that $\beta(1+r) = 1$. I discuss later the generalization to the case with $\beta(1+r) \neq 1$ and with 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—in the empirical 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 $e^{\eta_t^i}$. This means that $e^{p_t^i}$ is not just persistent but permanent, and its innovation multiplies future earnings at each period in the future by the same value it multiplies current earnings. In contrast, the realization of $e^{\varepsilon_t^i}$ only affects earnings at t. To obtain that people face a strictly positive amount of uncertainty, I assume $var(\varepsilon) > 0$. I let the variance of the persistent shocks be possibly equal to zero: $var(\eta) \geq 0$. 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 shock 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 constraint and positive spending constraint 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.

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,$e^{p_t^i}$, and $e^{\varepsilon_t^i}$:

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

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

2.2 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, however, 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, persistent earnings, and transitory earnings, except that from period t on their future earnings are 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] Note that, if people faced borrowing constraints, I would assume away those constraints under perfect foresight, 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. Precautionary saving, denoted PS_t, is the difference between

[5]For 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.

[6]The 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 $t+1$.

consumption under perfect foresight at t and consumption:

$$PS_t = c_t^{PF_t} - c_t \tag{2.7}$$

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:

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

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 parameters of the maximization problem are known and people can chose in perfect awareness their whole lifetime consumption path. The expected marginal utility of future consumption is the marginal utility of future consumption, so to equalize their marginal utility over time people equalize their consumption expenditures over time:

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

$$c_t^{PF} = c_{t+1}^{PF_t}. \tag{2.10}$$

There is no consumption growth in this simple model where $beta = (1+r)$. 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 expected resources (the one that would arise with quadratic preferences):

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

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. $\frac{1}{l_{t,T}} \to \frac{r}{1+r}$ when $T \to \infty$. 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. $u'''(.) > 0$ which is true of isoelastic utility functions), the presence of uncertainty

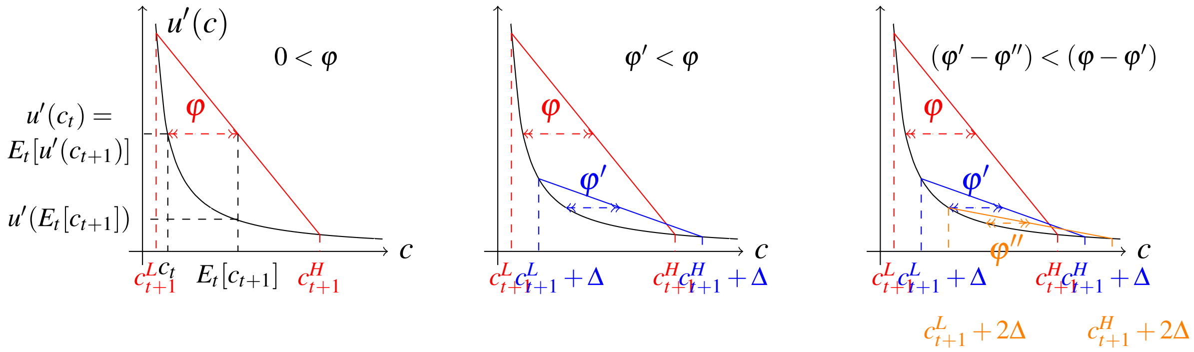

Figure 1: The precautionary premium over different intervals of future consumption

raises the expected marginal utility of future consumption above the marginal utility 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]

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

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

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 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 1 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^L_{t+1} and c^H_{t+1} 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

[7] Kimball 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).

$E_t[u'(c_{t+1})]$. The exact amount by which c_t must fall is φ_t such that $u'(c_t) = u'(E_t[c_{t+1}] - \varphi_t)$. I iterate forward, and plug the expressions of expected consumption growth at future period in the expected value of the intertemporal budget constraint. Iterating this reasoning forward, 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:

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

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 have the same effect on consumption 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 () with respect to a change in wealth:

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

Thus, a change in wealth leads people to adjust their current consumption for two reasons: first, a change in wealth modifies people's total expected lifetime resources, which is the only reason why consumption would respond under perfect foresight; second, a change in wealth 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

[8] A change in the transitory component of earnings e^{ε_t} has the same properties as a change in asset a_{t-1} because both affect consumption only through their impact on cash in hand $(1+r)a_{t-1} + y_t$ but by construction their magnitudes are different: a one unit change in wealth changes cash-in-hand by $(1+r)$ unit, while a change in the transitory component of earnings 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 assets does not. I abstract from this dependency on e^{p_t} by defining the MPC as the partial effect of assets. 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.

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:

$$(i)\ \frac{\partial c_t}{\partial a_{t-1}} > \frac{\partial c_t^{PF_t}}{\partial a_{t-1}} \ \ (\text{so } \tfrac{\partial PS_t}{\partial a_{t-1}} < 0)\ \text{, when } \tfrac{u'''(c)}{-u''(c)} \text{ strictly decreasing in c,}$$

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

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

Intuition for Lemma (i). When $\frac{u'''(.)}{-u''(.)}$ is strictly decreasing (which is true of any isoelastic utility function), the convexity of $-u''(.)$ is more pronounced than that of $u'(.)$, that is, $-u''(.)$ writes as an increasing and convex function of $u'(.)$. This 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 $u'(.)$ to $u'(.) - \Delta \times (-u''(.))$; the concavity of this new function is less pronounced because a relatively more convex element $(-u''(.))$ has been removed. As the convexity lessens over high intervals, so does the need for precautionary saving, and the premium φ_t gets smaller. The middle graph in Figure 1 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 given only partial results.[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) an increase in wealth above the target reduces saving thus reduces expected wealth growth; (2) a decrease in wealth below the target raises saving thus raises expected wealth growth; (3) expected wealth growth is zero at the target; combining the three, expected wealth growth is negative above and positive below the target.

Now, 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 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; conversely, if wealth increases, the decision problem shifts up and the level of saving decreases. 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 ensures 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 $u'''(.)u'(.) = k(-u''(.))^2$ with $k \neq 0$ (which

[9] More 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 $u'''(.)/(-u'')(.)$ 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.

[10] Kimball 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).

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 $u'''(.)$ is entirely determined by $u'(.)$ and $-u''(.)$: $u'''(.) = k(-u''(.))^2/u'(.)$. This ensures that $u'''(.)$ does not behave as $-u''(.)$ (for $k \neq 0$), and each shift up in the distribution of future consumption reduces the precautionary premium by less than the first shift when a shift up reduces the precautionary premium (and raises it by more than the first shift when a shift up raises the precautionary premium). The right graph in Figure 1 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, changes in wealth do 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 and show that it is negative. 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.

[11]The 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.

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

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

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

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

when $u(c) \in$ 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 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 decreasing relative prudence does the opposite. Note that the case with 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 new consumers' problem with one less state variable. Indeed, 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 homogeneity result tells us that these are irrelevant for the consumption solution.[12] Relatedly, the proof implies that it is also possible to normalize the consumers' problem by current wealth. Finally, the proof 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]

[12] I 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.

[13] The homogeneity result also clarifies the extent to which this model implies a linearity of consumption in his

Intuition for Lemma (iv). By Euler's homogeneous function theorem, the MPC function $(\partial c_t/\partial a_{t-1}) = (\partial c_t/\partial a_{t-1})(a_{t-1}, e^{p_t}, e^{\varepsilon_t})$, is homogeneous of degree zero in a_{t-1} and e^{p_t} if and only if $0 = a_{t-1}\frac{\partial^2 c_t}{\partial a_{t-1}^2} + e^{p_t}\frac{\partial^2 c_t}{\partial a_{t-1}e^{p_t}}$. 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 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 decreasing relative prudence does the opposite. The detailed proof is in Appendix A.5.

2.3 Persistent earnings and precautionary behavior

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

$$(i)\quad \frac{\partial c_t}{\partial e^{p_t}} < \frac{\partial c_t^{PF_t}}{\partial e^{p_t}} \quad (\text{so } \tfrac{\partial PS_t}{\partial e^{p_t}} > 0)$$

when $\frac{u'''(c)}{-u''(c)} > 0$ strictly decreasing in c & $\frac{u'''(c)c}{-u''(c)} > 1$ increasing in c,

$$(ii)\quad \frac{\partial^2 c_t}{\partial a_{t-1}e^{p_t}} > \frac{\partial^2 c_t^{PF_t}}{\partial a_{t-1}e^{p_t}} = 0 \quad (\text{so } \tfrac{\partial^2 PS_t}{\partial a_{t-1}e^{p_t}} < 0)$$

when $u(c) \in$ HARA & $\frac{u'''(c)c}{-u''(c)} > 1$ decreasing in c.

Graphical intuition. Figure 2 presents the mechanism graphically. I assume that the utility is isoelastic, which means that $\frac{u'''(c)}{-u''(c)} > 0$ is strictly decreasing in c, $\frac{u'''(c)c}{-u''(c)} > 1$ is constant, and utility is HARA, so all four Lemmas hold. The plain black line plots the evolution of consumption c_t with wealth a_{t-1} at a given level of persistent earnings e^{p_t}. 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. 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 a convex marginal utility make positive precautionary savings).

As implied by Lemma (i), 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).[14] This

time-invariant individual definition of permanent earnings: 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, when a_{t-1} is not proportional to e^{p_t} (which happens when people experience shocks to their permanent earnings that move them away from proportionality with wealth, or when people face no shocks to permanent earnings but their initial assets a_0 are not 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}.

[14]In fact, the slope of the plain back line at given of wealth is steeper than the slope of the dashed line at any level of wealth, since the slope of the dashed line is constant.

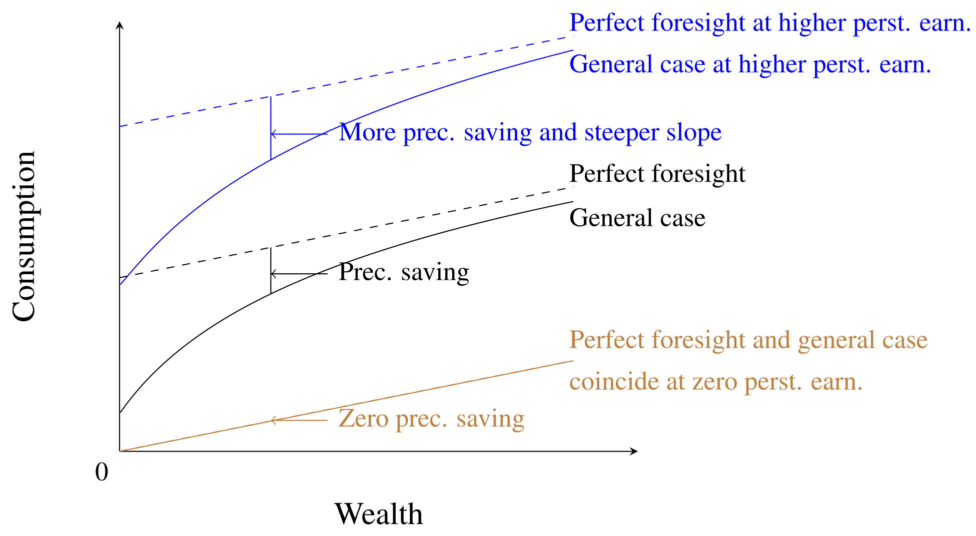

Figure 2: The evolution of consumption with wealth at different levels of persistent earnings

is because, 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). As implied by Lemma (ii), the plain line increases concavely with wealth. That is because 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 intuition of the Theorem is that the effect of precautionary saving is generally enhanced at higher levels of persistent earnings. This because, persistent earnings multiplies the risky part of people's revenue: when consumption is homogeneous in wealth and in persistent earnings, an increase in persistent earnings is equivalent to scaling up the decision problem but reducing wealth. I plot in brown the limit situation in which persistent earnings approaches zero ($p_t \to -\infty$ thus $e^{p_t} \to 0$). In that case, current and future expected earnings are all zero and the consumer can only count on its wealth to finance its current and future consumption. However, there is no longer any uncertainty: people are sure that, regardless of the shocks that realize, their future income will be zero. That is why the consumption function coincides with its perfect foresight counterpart and there is no precautionary saving. The slope of the consumption function is less steep than in the baseline (black) situation and equal to its perfect foresight value: an increase in assets no longer reduce the need for precautionary saving because people do not make any precautionary saving in that situation. In contrast, I plot in blue a situation in which persistent earnings is higher than in the baseline (black) situation. 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 in 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.

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

$$\frac{\partial c_t}{\partial e^{p_t}} \leq \frac{c_t}{e^{p_t}} - \frac{a_{t-1}}{e^{p_t}}\frac{\partial c_t}{\partial a_{t-1}} < \frac{c_t^{PF_t}}{e^{p_t}} - \frac{a_{t-1}}{e^{p_t}}\frac{\partial c_t^{PF_t}}{\partial a_{t-1}} = \frac{\partial c_t^{PF_t}}{\partial e^{p_t}}. \tag{2.16}$$

I move from the second to the third expression using that $\frac{c_t}{e^{p_t}} < \frac{c_t^{PF_t}}{e^{p_t}}$ because $c_t < c_t^{PF_t}$ when $\frac{u'''(c)}{-u''(c)} > 0$, that $a_{t-1} > 0$, and that $\frac{\partial c_t}{\partial a_{t-1}} > \frac{\partial c_t}{\partial a_{t-1}}$ when $\frac{u'''(c)}{-u''(c)} > 0$ is strictly decreasing in c (Lemma (i)). 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 did not actually multiply 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 it does not multiply. 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 decreasing in c, 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:

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

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

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 multiply by $(1+1/e^{p_t})$. The scaling up does not affect the MPC, but the reduction in wealth raises it when consumption is concave in wealth (because utility is HARA). The MPC is therefore higher at higher levels of persistent earnings because wealth is relatively smaller at higher levels of persistent earnings.

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 $V_t(a_{t-1}, e^{p_t}, e^{\varepsilon_t}) = \max_c \ u(c) + \beta_{t+1} e^{\delta z_{t+1} - z_t} E_t \left[V_{t+1}(a_t, e^{p_{t+1}}, e^{\varepsilon_{t+1}})\right]$ and the period budget constraint is $a_t^i = (1+r_t)a_{t-1} + y_t - c$, the Euler equation becomes:

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

with $R_{t,t+s} \equiv \Pi_{k=1}^{s} \beta_{t+k}(1+r_{t+k-1}) e^{\delta(z_{t+k} - z_{t+k-1})}$ 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 $l_{t,T} = \sum_{s=0}^{T-t} \frac{R_{t,t+s}^{1/\rho}}{\Pi_{k=0}^{s}(1+r_{t+k})}$. The Lemma (ii), (iii) and (iv) are unaffected by the presence of the additional term $R_{t,t+1}$, as discussed in each of their proofs, so Theorem (ii) is unaffected as well. The Lemma (i) still holds when the function $g(.) = (-u'') \circ (u')^{-1}(.)$ is such that $\frac{g'(yR^{-1})R}{g'(y)R} - \frac{g(yR^{-1})}{g(y)} \geq 0$, as in discussed in its proof. The isoelastic utility function and a range of exponential utility functions verify this additional condition, so Theorem (i) still holds under this additional conditions.

General conditions The mechanism at play in the standard life-cycle model might hold true more generally, when the persistent component of earnings shifts the precautionary motive not juts 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 2 might equivalently be caused by these other channels.

I thus generalize the framework to situation in which people decide on their consumption, which they finance out of N different resources, without specifying the exact problem they solve, nor the exact shape of the earnings process they face. I then exhibit in Appendix A.6

three proximate conditions that are sufficient to ensure that Theorem (i) holds, and three that are sufficient to ensure that Theorem (ii) hold. These conditions impose restrictions on the shape of consumption absent any uncertainty, restrictions on the effect of the resources that are not persistent earnings on precautionary saving and the precautionary component of the MPC, and restrictions on the homogeneity of consumption and of the MPC in persistent earnings and in the other resources.

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.[15]

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

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

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 select which of 14 possible categories of non-housing wealth their household belongs to. I check in Appendix B.3 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 loose 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 assumption 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).[16] I present descriptive statistics of my main variables in Appendix B.4.

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 administrative US data on earnings well.[17] It is:

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

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

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

This expression states that annual earnings are the product of a dummy for employment status ν_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, ν_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

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

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

one year). The probability to be nonemployed at t, $p_{\nu_{t-1}^i}$ 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 $p_{\nu_{t-1}^i}$ 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, 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 $perst_t^i$

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

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.[18]

[18]Formally, the average rescaled transitory component in the population is $E[e^{\varepsilon_t^i - \overline{\varepsilon}}] > e^{E[\varepsilon_t^i] - \overline{\varepsilon}} = 1$, the average rescaled fixed effect is $E[e^{\alpha^i - \overline{\alpha}}] > e^{E[\alpha^i] - \overline{\alpha}} = 1$, and the average rescaled age trend is $E[e^{g(t) - \overline{g(t)}}] > e^{E[g(t)] - \overline{g(t)}} = 1$. 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.

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 $t+1$

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

Dividing expected future annual earnings by the probability to be employed at the next period and taking the log of the resulting term yields

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

Thus, among employed people, the residual from a regression of $ln(E_t^i[y_{t+1}]/(1-p_{v_t}^i))$ 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.[19] 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 $perst_t^i = e^{p_t^i} e^{\overline{\varepsilon}} e^{\overline{\alpha}} e^{\overline{g(t)}}$:

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

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

[19] It 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.

overoptimistic.[20]

	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

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 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.5 that the method does not rely on

[20]Also 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.

any specific assumption about the timing of the persistent and transitory shocks within the year (and shocks occurring within the year can be consistent with the distributions I later assume in numerical simulations as discussed in B.6). I also explain in Appendix B.5 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.5 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 $perst_t^i$ but $perst_t^i e^{\varepsilon_{t+1}/\rho}$, which could bias the results. I show in Appendix B.7 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.8 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, 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 in Appendix B.9 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, 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:

$$\begin{aligned} MPC_t^i = a_1 + a_2\ perst_t^i(1 + b_2 hh\ size_t^i) + a_3\ earn_t^i(1 + b_3 hh\ size_t^i) + a_4\ wealth\ cat_t^i(1 + b_4 hh\ size_t^i) \\ + a_5\ hh\ size_t^i + a_6 state_t^i + a_7\ fixed\ dem^i + a_8 period_t + a_9 age\ group_t^i + \xi_t^i, \end{aligned} \tag{4.1}$$

with MPC_t^i denoting the reported MPC out of hypothetical shocks of respondent i at period t, $pers_t^i$ 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 $perst_t^i$ is built as described in the previous subsection.

	MPC neg.	MPC pos.
Persistent earnings in \$10,000 (a_2)	0.015***	0.014**
	(0.005)	(0.006)
Earnings in \$10,000 (a_3)	-0.009**	-0.009
	(0.005)	(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 16 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 17 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 18 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 $\frac{e^{p_t^i}}{1+e^{p_t^i}} < 1$— 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 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 19 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 20 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 21 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 22 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 18 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 I 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

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

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 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, as reported in Table 27 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.[21] 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

[21]For 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.

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 higher[22]—, and (ii) that narrower targeting would also means 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 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: including the nonemployed, the average MPCs out of negative and positive shocks in 10th income percentile group are 0.84 and 0.58 while in the whole sample they are 0.79 and 0.55.

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

[22]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.

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 $>=$ 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.[23] 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 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 $u(.)$ 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

[23]More 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).

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.[24]

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.[25] 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 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.[26]

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.[27] 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 $\beta = 0.951$.

[24] See Figure 1 in Attanasio, Banks, Meghir, and Weber 1999 and Figure 4 in Attanasio 1999.

[25] The 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.

[26] Incidentally, 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.

[27] This 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. The questions I use are 'Approximately what is the total current value of your [and your spouse's/partner's] savings and investments (such as checking and savings accounts, CDs, stocks, bonds, mutual funds, Treasury bonds), excluding those in retirement accounts?' (D16new) for liquid wealth and 'Next consider all outstanding debt you [and your spouse/partner] have, including balances on credit cards (including retail cards), auto loans, student loans, other personal loans, as well as medical or legal bills, but excluding all housing related debt (such as mortgages, home equity lines of credit, home equity loans). Approximately, what is the total amount of your [and your spouse's/partner's] current outstanding debt?'. 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.

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 than on its previous value:

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

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

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

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

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

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

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

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

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, $tax(y_t^i) = \tau^b \big(y_t^i - ((y_t^i)^{-\tau^\rho} + \tau^s)^{-1/\tau^\rho}\big)$ parametrized with $\tau^b = 258$, $\tau^\rho = 0.768$, $\tau^s = 2.0e-4$ as in Kaplan and Violante 2010.[28]

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.[29]

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

[28]Contrary to Kaplan and Violante 2010 who model net income and use the inverse of the tax function to recover gross income, here, what I model is pre-tax earnings and I use the tax function to recover net earnings. People pay 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.

[29]The 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.

deflated and expressed in 2010$ value. For comparability with my survey data, I convert my simulated data from 2010$ to 2014$.[30]

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.

	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

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.

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 +/-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.

[30]This 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.

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:

$$MPC_t^i = a_1 + a_2\ perst_t^i + a_3\ earn_t^i + a_4\ wealth\ cat_t^i + a_7 t + \xi_t^i. \tag{5.9}$$

	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%).

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

	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

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 at the same value as in the baseline model, 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 than in the baseline one, 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 is the same as in the simulations of the baseline model. 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 is the same as in the simulations of the baseline model. 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 the effect of persistent earnings conditional on wealth and persistent earnings is stronger, or because the distributions of wealth and persistent earnings have changed in a way that raises the average effect. Adding the variance dependency also increases slightly the magnitude of the effect of total earnings (conditional on persistent earnings), and makes its point estimates negative.

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. Adding the redistribution system also reduces the effect of persistent earnings on the MPC out of a positive shock. 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 although 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's MPC decreases with the persistent component of their earnings, 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., Blundell, R., and Bonhomme, S. (2017). “Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework.” *Econometrica* 85, pp. 693–734.

Arellano, Manuel, Bonhomme, Stéphane, Vera, Micole De, Hospido, Laura, and Wei, Siqi (Sept. 2021). *Income Risk Inequality: Evidence from Spanish Administrative Records*. Working Papers 2136. Banco de España.

Armantier, Olivier, Topa, Giorgio, Klaauw, Wilbert Van der, and Zafar, Basit (2017). “An overview of the Survey of Consumer Expectations.” *Economic Policy Review* 23-2, pp. 51–72.

Arrow, K. (1965). “Aspects of the theory of risk-bearing.” Yrjö Jahnssonin Säätiö (Yrjo Jahnsson Lectures).

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, Adrien (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, Almut, Duernecker, Georg, Forstner, Susanne K., and Goensch, Johannes (2021). *The Effects of Biased Labor Market Expectations on Consumption, Wealth Inequality, and Welfare*. CESifo Working Paper Series 9326. CESifo.

Blundell, R., Pistaferri, L., and Preston, I. (2008). “Consumption Inequality and Partial Insurance.” *American Economic Review* 98.5, pp. 1887–1921.

Borusyak, Kirill, Jaravel, Xavier, and Spiess, Jann (Apr. 2022). *Revisiting Event Study Designs: Robust and Efficient Estimation*. CEPR Discussion Papers 17247. C.E.P.R. Discussion Papers. URL: `https://ideas.repec.org/p/cpr/ceprdp/17247.html`.

Boutros, M. (2021). “Household Finances and Fiscal Stimulus in 2008.” *mimeo*.

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

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

Bunn, Philip, Le Roux, Jeanne, Reinold, Kate, and Surico, Paolo (2018). "The consumption response to positive and negative income shocks." *Journal of Monetary Economics* 96, pp. 1–15. ISSN: 0304-3932.

Carroll, C. (1997). "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis." *The Quarterly Journal of Economics* 112.1, pp. 1–55.

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.

Carroll, Christopher (2006). "The method of endogenous gridpoints for solving dynamic stochastic optimization problems." *Economics Letters* 91.3, pp. 312–320. URL: `https://EconPapers.repec.org/RePEc:eee:ecolet:v:91:y:2006:i:3:p:312-320`.

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

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.

Felgueroso, Florentino, García-Pérez, José-Ignacio, Jansen, Marcel, and Troncoso-Ponce, David (2018). "The Surge in Short-Duration Contracts in Spain." *De Economist* 166.4, pp. 503–534.

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.

Gelman, M., Kariv, S, Shapiro, M., Silverman, D., and Tadelis, S (2018). "How individuals respond to a liquidity shock: Evidence from the 2013 government shutdown." *Journal of Public Economics*.

Golosov, Mikhail, Graber, Michael, Mogstad, Magne, and Novgorodsky, David (2021). *How Americans Respond to Idiosyncratic and Exogenous Changes in Household Wealth and Unearned Income*. Working Paper 29000. National Bureau of Economic Research.

Gong, Liutang, Zhong, Ruquan, and Zou, Heng-fu (2012). "On the concavity of the consumption function with the time varying discount rate." *Economics Letters* 117.1, pp. 99–101.

Gouveia, Miguel and Strauss, Robert P. (1994). "Effective Federal Individual Tax Functions: An Exploratory Empirical Analysis." *National Tax Journal* 47.2, pp. 317–39.

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

Hurd, M. and Rohwedder, S. (2013). "Heterogeneity in spending change at retirement." *The Journal of the Economics of Ageing* 1, pp. 60–71.

Jappelli, Tullio and Pistaferri, Luigi (2014). "Fiscal Policy and MPC Heterogeneity." *American Economic Journal: Macroeconomics* 6.4, pp. 107–36.

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., Violante, G., and Weidner, J. (2014). "The Wealthy Hand-to-Mouth." *Brookings Papers on Economic Activity*. ISSN: 1533-4465.

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

Karger, Ezra and Rajan, Aastha (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, Daniel J., Melcangi, Davide, and Pilossoph, Laura (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.

Orchard, Jacob, Ramey, Valerie A., and Wieland, Johannes (2022). "Micro MPCs and Macro Counterfactuals: The Case of the 2008 Rebates." *Working paper*.

Palmer, James A. (2003). *Relative Convexity*. Tech. rep. ECE Dept., UCSD.

Parker, J., Souleles, N., Johnson, D., and McClelland, R. (2013). "Consumer Spending and the Economic Stimulus Payments of 2008." *American Economic Review* 103.6, pp. 2530–53.

Parker, Jonathan A, Schild, Jake, Erhard, Laura, and Johnson, David (2022). *Household Spending Responses to the Economic Impact Payments of 2020: Evidence from the Consumer Expenditure Survey*. Working Paper 29648. National Bureau of Economic Research.

Pistaferri, Luigi (2001). "Superior Information, Income Shocks, And The Permanent Income Hypothesis." *The Review of Economics and Statistics* 83.3, pp. 465–476.

Pratt, J. (1964). "Risk Aversion in the Small and in the Large." *Econometrica* 32.1/2, pp. 122–136.

Rozsypal, Filip and Schlafmann, Kathrin (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, Alexis Akira (2021). "Necessity of hyperbolic absolute risk aversion for the concavity of consumption functions." *Journal of Mathematical Economics* 94.C.

A Theoretical results

A.1 The expression of precautionary saving in a multiperiod model

In the multiperiod model described by (2.1)-(2.6) the precautionary premium φ measures the effect of uncertainty on expected consumption growth between two periods, not the effect of uncertainty on the level of consumption, that is, not precautionary saving. To get an expression of precautionary saving in a multiperiod model, I iterate forward the reasoning by which I express expected consumption growth between t and $t+1$. I obtain that $c_t = E_t[c_{t+s}] - \sum_{j=1}^{s} E_t[\varphi_{t+j-1}]$ for all $1 \leq s \leq T-t$: current consumption equals expected future consumption at $t+s$, minus some precautionary consumption growth that writes as a sum of the expected precautionary premia between consecutive periods. Precautionary saving eventually writes as a double sum of current and expected precautionary premia between consecutive periods. To show this, I substitute each $E_t[c_{t+s}]$ with $c_t + \sum_{j=1}^{s} E_t[\varphi_{t+j-1}]$ for all $1 \leq s \leq T-t$ in the expected intertemporal budget constraint and rearrange to obtain the following equilibrium relationship for consumption:

$$\begin{aligned}
(1+r)a_{t-1} + \sum_{s=0}^{T-t} \frac{E_t[y_{t+s}]}{(1+r)^s} &= \sum_{s=0}^{T-t} \frac{E_t[c_{t+s}]}{(1+r)^s} \qquad \text{(A.1)}\\
&= c_t \left(\sum_{s=0}^{T-t} \frac{1}{(1+r)^s} \right) + \sum_{s=1}^{T-t} \sum_{j=1}^{s} E_t[\varphi_{t+j-1}] \frac{1}{(1+r)^s} \\
&= c_t \left(\sum_{s=0}^{T-t} \frac{1}{(1+r)^s} \right) + \sum_{j=1}^{T-t} \frac{E_t[\varphi_{t+j-1}]}{(1+r)^j} \sum_{s=j}^{T-t} \frac{1}{(1+r)^{s-j}} \\
&= c_t \underbrace{\left(\sum_{s=0}^{T-t} \frac{1}{(1+r)^s} \right)}_{l_{t,T}} + \sum_{j=1}^{T-t} \frac{E_t[\varphi_{t+j-1}]}{(1+r)^j} \underbrace{\sum_{s=0}^{T-(t+j)} \frac{1}{(1+r)^s}}_{l_{t+j,T}} .
\end{aligned}$$

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. $\frac{1}{l_{t,T}} \to \frac{r}{1+r}$ when $T \to \infty$. This eventually yields the following expression of consumption:

$$c_t = \underbrace{\frac{1}{l_{t,T}} \left((1+r)a_{t-1} + \sum_{s=0}^{T-t} \frac{E_t[y_{t+s}]}{(1+r)^s} \right)}_{c_t^{PF_t}} - \underbrace{\frac{1}{l_{t,T}} \left(\sum_{j=1}^{T-t} l_{t+j,T} \frac{E_t[\varphi_{t+j-1}]}{(1+r)^j} \right)}_{PS_t} \qquad \text{(A.2)}$$

Note that this is not a closed-form expression of consumption but an equilibrium condition, because the equity premia are determined jointly with c_t. This expression indicates that, under perfect foresight, people would consume the exogenous fraction $1/l_{t,T}$ of their lifetime expected resources. In the presence of uncertainty, however, people set their current consumption below their expected future consumption and expect to do so in the future, so they net out from their lifetime expected resources the precautionary consumption growth that they plan to implement at each future period. They consume a fraction $1/l_{t,T}$ of what remains instead of a fraction of their full resources. The share of total expected resources optimally allocated to the current period is no longer exogenous and equal to $1/l_{t,T}$ because there are extra transfers of resources from the current to the future periods captured by the precautionary saving term.

A.2 Proof of Lemma (i)

The proof of Lemma (i) that $\frac{\partial c_t}{\partial a_{t-1}} > \frac{\partial c_t^{PF_t}}{\partial a_{t-1}}$ is by backward induction. First, I establish that a non-strict version of Lemma (i) is true at the last period T. Second, I establish that if a non-strict version of Lemma (i) is true at $t+1$, then Lemma (i) is strictly true at t. This establishes that Lemma (i) is strictly true at all periods $t < T$.

At $t = T$, $c_T = (1+r)a_{T-1} + y_T = c_T^{PF_T}$ so the MPC is the same in the presence of uncertainty as under perfect foresight at T (or at any perfect foresight horizon):

$$\frac{\partial c_T}{\partial a_{T-1}} = 1 + r = \frac{\partial c_T^{PF_T}}{\partial a_{T-1}}.$$

Thus, a non-strict version of Lemma (i) is true at T.

Now I assume that a non-strict version of Lemma (i) is true at $t+1$. I differentiate both sides of the Euler equation (2.8) with respect to a change in a_{t-1}, divide both sides by $(-u''(c_t))$, and rearrange:

$$\begin{aligned}
\frac{\partial c_t}{\partial a_{t-1}} &= E_t[\frac{\partial c_{t+1}}{\partial a_{t-1}} \frac{-u''(c_{t+1})}{-u''(c_t)}] \\
\frac{\partial c_t}{\partial a_{t-1}} &= E_t[\frac{\partial a_t}{\partial a_{t-1}} \frac{\partial c_{t+1}}{\partial a_t} \frac{-u''(c_{t+1})}{-u''(c_t)}] \\
\frac{\partial c_t}{\partial a_{t-1}} &= E_t[((1+r) - \frac{\partial c_t}{\partial a_{t-1}}) \frac{\partial c_{t+1}}{\partial a_{t-1}} \frac{-u''(c_{t+1})}{-u''(c_t)}] \\
\frac{\partial c_t}{\partial a_{t-1}} &= (1+r) \frac{E_t[\frac{\partial c_{t+1}}{\partial a_t} \frac{-u''(c_{t+1})}{-u''(c_t)}]}{1 + E_t[\frac{\partial c_{t+1}}{\partial a_t} \frac{-u''(c_{t+1})}{-u''(c_t)}]} = (1+r) g\left(E_t[\frac{\partial c_{t+1}}{\partial a_t} \frac{-u''(c_{t+1})}{-u''(c_t)}] \right),
\end{aligned} \tag{A.3}$$

with $g(x) = \frac{x}{1+x}$. This function is strictly increasing over $[0, +\infty[$. It is also positive and strictly below one over this interval, so $\frac{\partial c_t}{\partial a_{t-1}}$ is positive and strictly below $(1+r)$. Now, by assumption,

$\frac{\partial c_{t+1}}{\partial a_t} \geq \frac{\partial c_{t+1}^{PF_{t+1}}}{\partial a_t}$. I examine what it implies for $\frac{\partial c_t}{\partial a_{t-1}}$:

$$\frac{\partial c_t}{\partial a_{t-1}} = (1+r)g\left(E_t[\frac{\partial c_{t+1}}{\partial a_t}\frac{-u''(c_{t+1})}{-u''(c_t)}]\right) \tag{A.4}$$

$$\geq (1+r)g\left(E_t[\frac{\partial c_{t+1}^{PF_{t+1}}}{\partial a_t}\frac{-u''(c_{t+1})}{-u''(c_t)}]\right) \tag{A.5}$$

$$\geq (1+r)g\left(\frac{1+r}{l_{t+1,T}}E_t[\frac{-u''(c_{t+1})}{-u''(c_t)}]\right) \tag{A.6}$$

$$> (1+r)g\left(\frac{1+r}{l_{t+1,T}}\right) = \frac{1+r}{1+l_{t+1,T}/(1+r)} = \frac{1+r}{l_{t,T}} = \frac{\partial c_t^{PF_t}}{\partial a_{t-1}} \tag{A.7}$$

I move from the first line to the second using that a non-strict version of Lemma (i) is true at $t+1$. I move from the second to the third line using that, under perfect foresight, $\frac{\partial c_{t+1}}{\partial a_t} = \frac{1+r}{l_{t+1,T}}$, from expression (2.11). I move from the third to the fourth line using that, because $\frac{u'''(.)}{-u''(.)}$ is strictly decreasing, the absolute risk aversion associated with $-u''(.)$ is strictly larger than the absolute risk aversion associated with $u'(.)$ at every level of consumption. Indeed, $\frac{u'''(.)}{-u''(.)}$ being strictly decreasing is equivalent to $\frac{-u''''(.)}{u'''(.)} > \frac{u'''(.)}{-u''(.)}$, and $r_1(.) = \frac{-u''''(.)}{u'''(.)}$ is the risk aversion associated with $u_1(.) = -u''(.)$ while $r_2(.) = \frac{u'''(.)}{-u''(.)}$ is the risk aversion associated with $u_2(.) = u'(.)$. From Pratt 1964-Arrow 1965, this means that the premium φ_t^{r1} associated with $-u''(.)$ to compensate for the risk in the distribution of future consumption strictly larger than the premium $\varphi_t^{r2} = \varphi_t$ associated with $u'(.)$ to compensate for the same risk. Therefore:

$$E_t[-u''(c_{t+1})] = -u''(E_t[c_{t+1}] - \varphi_t^{r2}) > -u''(E_t[c_{t+1}] - \varphi_t) = -u''(c_t),$$

and the ratio $E_t[-u''(c_{t+1})]/(-u''(c_t))$ is strictly larger than one, thus strictly larger than its value under perfect foresight at t. Finally, in the fourth line, I use that $g(1/y) = (1/y)/(1+1/y) = 1/(1+y)$ and that by construction $1 + l_{t+1,T}/(1+r) = l_{t,T}$.

In the intuition, I use the notion that $-u''(.)$ is a convex transformation of $u'(.)$ because I find it easy to relate to but it directly maps into the proof above. Indeed, Palmer 2003's Theorem 4 establishes that a given twice differentiable function u_1 is strictly convex relative to another twice differentiable function u_2, that is, that there exists a strictly increasing and strictly convex function $h(.)$ such that $u_1(.) = h(u_2(.))$, if and only if $\frac{u_1''}{|u_1'|} > \frac{u_2''}{|u_2'|}$, which holds true for $u_1(.) = -u''(.)$ and $u_2(.) = u'(.)$ when $\frac{u'''(.)}{-u''(.)}$ is strictly decreasing. This means that, when $u'''(.)/-u''(.)$ is strictly decreasing, $-u''(.)$ is strictly convex relative to $u'(.)$ and the premium associated with $-u''(.)$ is larger than the premium associated with $u'(.)$ because a shift up in the consumption problem by a small amount Δ is equivalent to a shift in the marginal utility from $u'(.)$ to $u'(.) - \Delta(-u''(.))$, which is relatively less convex.[31]

[31] The fact that $u'(.) - \Delta \times (-u''(.))$ writes as a strictly increasing and concave function of $u'(.)$ is an application of Palmer 2003's Lemma 1 that if u_2 and u_3 are concave relative to u_1, then their sum is so as well. Here $u'(.)$ is concave (not strictly !) relative to $u'(.)$, and $-\Delta(-u''(.))$ is strictly concave relative to $u'(.)$ for any $\Delta > 0$ (because

Time-varying demographics, discount factors, and interest rates. In a model with time-varying demographic characteristics, discount factors and interest rates, the Euler equation becomes:

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

with $R_{t,t+s} \equiv \Pi_{k=1}^{s}\beta_{t+k}(1+r_{t+k-1})e^{\delta(z_{t+k}-z_{t+k-1})}$. Thus, an optimizing household equalizes its expected marginal utility over time, weighted by the strength of intertemporal substitution motives. The last period is unchanged. Under perfect foresight at t:

$$u'(c_t^{PF_t}) = u'(c_{t+1}^{PF_t})R_{t,t+1}. \tag{A.8}$$

I also differentiate both sides of (A.8) with respect to a change in a_{t-1} and rearrange:

$$\begin{aligned}
\frac{\partial c_t^{PF_t}}{\partial a_{t-1}} &= \frac{\partial c_{t+1}^{PF_t}}{\partial a_{t-1}}\frac{-u''(c_{t+1}^{PF_t})}{-u''(c_t^{PF_t})}\\
\frac{\partial c_t^{PF_t}}{\partial a_{t-1}} &= \frac{\partial a_t^{PF_t}}{\partial a_{t-1}}\frac{\partial c_{t+1}^{PF_t}}{\partial a_t^{PF_t}}\frac{-u''(c_{t+1}^{PF_t})}{-u''(c_t^{PF_t})}\\
\frac{\partial c_t^{PF_t}}{\partial a_{t-1}} &= \big((1+r)-\frac{\partial c_t^{PF_t}}{\partial a_{t-1}}\big)\frac{\partial c_{t+1}^{PF_t}}{\partial a_{t-1}}\frac{-u''(c_{t+1}^{PF_t})}{-u''(c_t^{PF_t})}\\
\frac{\partial c_t^{PF_t}}{\partial a_{t-1}} &= (1+r)g\left(\frac{\partial c_{t+1}^{PF_t}}{\partial a_t^{PF_t}}\frac{-u''(c_{t+1}^{PF_t})}{-u''(c_t^{PF_t})}R_{t,t+1}\right)\\
\frac{\partial c_t^{PF_t}}{\partial a_{t-1}} &= (1+r)g\left(\frac{\partial c_{t+1}^{PF_t}}{\partial a_t^{PF_t}}\frac{-u''\big((u')^{-1}(u'(c_t^{PF_t})R_{t,t+1}^{-1})\big)}{-u''(c_t^{PF_t})}R_{t,t+1}\right)
\end{aligned}$$

Considering the general case with uncertainty, i also differentiate both sides of (2.18) with respect to a change in a_{t-1}, and do a similar rearranging:

$$\frac{\partial c_t}{\partial a_{t-1}} = (1+r)g\left(E_t\big[\frac{\partial c_{t+1}}{\partial a_t}\frac{-u''(c_{t+1})}{-u''(c_t)}\big]R_{t,t+1}\right) \tag{A.9}$$

$$\geq (1+r)g\left(\frac{\partial c_{t+1}^{PF_{t+1}}}{\partial a_t}\frac{E_t\big[-u''(c_{t+1})\big]}{-u''(c_t)}R_{t,t+1}\right) \tag{A.10}$$

$$> (1+r)g\left(\frac{\partial c_{t+1}^{PF_t}}{\partial a_t^{PF_t}}\frac{-u''\big((u')^{-1}(u'(c_t)R_{t,t+1}^{-1})\big)R_{t,t+1}}{-u''(c_t)}\right) \tag{A.11}$$

$$> (1+r)g\left(\frac{\partial c_{t+1}^{PF_t}}{\partial a_t^{PF_t}}\frac{-u''\big((u')^{-1}(u'(c_t^{PF_t})R_{t,t+1}^{-1})\big)R_{t,t+1}}{-u''(c_t^{PF_t})}\right) = \frac{\partial c_t^{PF_t}}{\partial a_{t-1}} \tag{A.12}$$

I move from the first to the second line using that a non-strict version of Lemma (i) holds true

$(-u''(.))$ is strictly convex relative to $u'(.)$ so its opposite is strictly concave), so their sum is concave relative to $u'(.)$.

at $t+1$. I move from the second to the third line using that the MPC at $t+1$ is the same under perfect foresight at $t+1$ as under perfect foresight at t, because the MPC is exogenous and independent of the realization of the shocks: $\frac{\partial c_{t+1}^{PF_{t+1}}}{\partial a_t} = \frac{\partial c_{t+1}^{PF_t}}{\partial a_t^{PF_t}}$ (defining the functions $c_t^{PF_t} = c_t^{PF_t}(a_{t-1}, e^{p_t})$ and $c_{t+1}^{PF_t} = c_{t+1}^{PF_t}(a_t^{PF_t}, E_t[e^{p_{t+1}}])$). In addition, because the premium φ_t^{r1} associated with $(-u'')$ is strictly larger than the premium φ_t associated with $u'(.)$:

$$E_t\big[-u''(c_{t+1})\big] = -u''(E_t[c_{t+1}] - \varphi_t^{r1}) > -u''(E_t[c_{t+1}] - \varphi_t) = -u''\big((u')^{-1}(u'(c_t)R_{t,t+1}^{-1})\big),$$

with $E_t[c_{t+1}] - \varphi_t = (u')^{-1}(u'(c_t)R_{t,t+1}^{-1})$ from a rearranging of the Euler equation. Finally, I move from the third to the fourth line using that the expression $\frac{-u''\big((u')^{-1}(u'(c_t)R_{t,t+1}^{-1})\big)R_{t,t+1}}{-u''(c_t)}$ rewrites as $l(y,R) = \frac{g(yR^{-1})R}{g(y)}$, with $g(.) = (-u'')\circ(u')^{-1}(.)$, $y = u'(c_t)$, and $R = R_{t,t+1}$. This function $l(y,R)$ is constant or increasing in y when $g(.)$, y and R are such that $\frac{g'(yR^{-1})}{g'(y)(R^{-1})^{-1}} - \frac{g(yR^{-1})}{g(y)} \geq 0$.[32] When $l(y,R)$ is constant or increasing in y then $l(u'(c_t), R_{t,t+1}) \geq l(u'(c_t^{PF_t}), R_{t,t+1})$ because $c_t < c_t^{PF_t}$ and $u'(.)$ is strictly decreasing.

A.3 Proof of Lemma (ii)

The proof of Lemma (ii) that $\frac{\partial^2 c_t}{\partial a_{t-1}^2} > \frac{\partial^2 c_t^{PF_t}}{\partial a_{t-1}^2}$ is by backward induction as well. First, I establish that a non-strict version of Lemma (ii) is true at the last period T. Second, I establish that, if a non-strict version of Lemma (ii) is true at $t+1$, the it must be strictly true at t. This establishes the strict version of Lemma (ii) at all periods $t < T$.

At $t = T$, $c_T = (1+r)a_{T-1} + y_T$ so consumption is linear in wealth, both in the presence of uncertainty and under perfect foresight:

$$\frac{\partial^2 c_T}{\partial a_{T-1}^2} = 0 = \frac{\partial^2 c_T^{PF_t}}{\partial a_{T-1}^2}.$$

Thus, a non-strict version of Lemma (ii) is true at T.

Now I assume that a non-strict version of Lemma (ii) is true at $t+1$. I differentiate both

[32] Since $h(u'(c_t)) = (-u'')(c_t) > 0$ and $g'(u'(c_t)) = u'''(c_t)/(-u''(c_t) > 0$, the condition $\frac{g'(yR^{-1})}{g'(y)(R^{-1})^{-1}} - \frac{g(yR^{-1})}{g(y)} \geq 0$ is sufficient to ensure that $\frac{\partial l(y,R)}{\partial y} = \frac{g'(y)R}{g(y)}\big(\frac{g'(yR^{-1})}{g'(y)R} - \frac{g(yR^{-1})}{g(y)} \geq 0$. Note that this difference is zero when $g(.) = (-u'')\circ(u')^{-1}(.)$ is separable, that is, such that $g(yR^{-1}) = g_y(y)g_R(R^{-1})$. Indeed, then $(R^{-1})g'(yR^{-1}) = \frac{\partial g(yR^{-1})}{\partial y} = \frac{\partial g_y(y)g_R(R^{-1})}{\partial y} = g_y'(y)g_R(R^{-1})$ and $\frac{g'(yR^{-1})}{g'(y)(R^{-1})^{-1}} - \frac{g(yR^{-1})}{g(y)} = \frac{g_y'(y)g_R(R^{-1})(R^{-1})^{-1}}{g_y'(y)g_R(1)(R^{-1})^{-1}} - \frac{g_y(y)g_R(R^{-1})}{g_y(y)g_R(1)} = 0.$

sides of the Euler equation (2.8) twice with respect to a_{t-1} and rearrange:

$$\begin{aligned}
&\frac{\partial^2 c_t}{\partial a_{t-1}^2}(-u''(c_t)) - \big(\frac{\partial c_t}{\partial a_{t-1}}\big)^2 u'''(c_t) = \\
&\qquad E_t\big[\big(\frac{\partial^2 a_t}{\partial a_{t-1}^2}\frac{\partial c_{t+1}}{\partial a_t} + \big(\frac{\partial a_t}{\partial a_{t-1}}\big)^2\frac{\partial^2 c_{t+1}}{\partial a_t^2}\big)(-u''(c_{t+1}))\big] - E_t\big[\big(\frac{\partial c_{t+1}}{\partial a_{t-1}}\big)^2 u'''(c_{t+1})\big] \\
&\frac{\partial^2 c_t}{\partial a_{t-1}^2} = E_t\big[\big(\frac{\partial^2 a_t}{\partial a_{t-1}^2}\frac{\partial c_{t+1}}{\partial a_t} + \big(\frac{\partial a_t}{\partial a_{t-1}}\big)^2\frac{\partial^2 c_{t+1}}{\partial a_t^2}\big)\frac{-u''(c_{t+1})}{-u''(c_t)}\big] \\
&\qquad + \big(\frac{\partial c_t}{\partial a_{t-1}}\big)^2\frac{u'''(c_t)}{-u''(c_t)} - E_t\big[\big(\frac{\partial c_{t+1}}{\partial a_{t-1}}\big)^2\frac{u'''(c_{t+1})}{-u''(c_t)}\big] \\
&\frac{\partial^2 c_t}{\partial a_{t-1}^2}\left(1 + E_t\big[\frac{\partial c_{t+1}}{\partial a_t}\frac{-u''(c_{t+1})}{-u''(c_t)}\big]\right) = E_t\big[\big(\frac{\partial a_t}{\partial a_{t-1}}\big)^2\frac{\partial^2 c_{t+1}}{\partial a_t^2}\frac{-u''(c_{t+1})}{-u''(c_t)}\big] \\
&\qquad - \frac{u'''(c_t)}{-u''(c_t)}\left(E_t\big[\big(\frac{\partial c_{t+1}}{\partial a_{t-1}}\big)^2\frac{u'''(c_{t+1})}{u'''(c_t)}\big] - E_t\big[\frac{\partial c_{t+1}}{\partial a_{t-1}}\frac{-u''(c_{t+1})}{-u''(c_t)}\big]^2\right). \qquad \text{(A.13)}
\end{aligned}$$

I move from the first to the second line by dividing all terms by $(-u''(c_t))$, and shifting the second term on the left hand-side to the right hand-side. I move from the second to the third line by substituting $\frac{\partial^2 a_t}{\partial a_{t-1}^2}$ with $-(1+r)\frac{\partial^2 c_t}{\partial a_{t-1}^2}$ (since $a_t = (1+r)a_{t-1} + y_t - c_t$) and $\frac{\partial c_t}{\partial a_{t-1}} = E_t[\frac{\partial c_{t+1}}{\partial a_{t-1}}\frac{-u''(c_{t+1})}{-u''(c_t)}]$, from the first differentiation of the Euler equation with respect to a_{t-1}). I also shift all terms with $\frac{\partial c_t^2}{\partial a_{t-1}^2}$ on the left hand-side and factorize the last two terms of the right hand-side by $-\frac{u'''(c_t)}{-u''(c_t)}$.

In this last expression (A.13), I now use that, when utility is HARA, $u(.)$ is such that $u'''(.) = k(-u''(.))^2/u'(.)$ (with $k \neq 0$ when utility is non-quadratic as is the case here because $u'''(.) > 0$), to substitute $\frac{u'''(c_{t+1})}{u'''(c_t)}$ with $\frac{(-u''(c_{t+1}))^2/u'(c_{t+1})}{(-u''(c_t))^2/u'(c_t)}$. I can then rewrite (A.13) using the more compact notations $A_{t+1} = \frac{\partial c_{t+1}}{\partial a_t}\frac{-u''(c_{t+1})}{-u''(c_t)}$, $B_{t+1} = \frac{\partial c_{t+1}}{\partial a_{t-1}}\frac{-u''(c_{t+1})}{-u''(c_t)}$, and $U_{t+1} = \frac{u'(c_{t+1})}{u'(c_t)}$:

$$\frac{\partial^2 c_t}{\partial a_{t-1}^2} = \underbrace{\frac{1}{1+E_t[A_{t+1}]}}_{>0}\left(E_t\big[\big(\frac{\partial a_t}{\partial a_{t-1}}\big)^2\underbrace{\frac{\partial^2 c_{t+1}}{\partial a_t^2}}_{\leq 0}\frac{-u''(c_{t+1})}{-u''(c_t)}\big] - \frac{u'''(c_t)}{-u''(c_t)}\underbrace{\left(E_t[B_{t+1}^2 U_{t+1}^{-1}] - E_t[B_{t+1}]^2\right)}_{>0 \text{ with Cauchy-Schwartz}}\right) < 0$$

The term $E_t[A_{t+1}]$ is strictly positive because the terms that compose A_{t+1} are. The term $\frac{\partial^2 c_{t+1}}{\partial a_t^2}$ is negative (not necessarily strictly) from the assumption that a non-strict version of Lemma (ii) is true at $t+1$. Finally, the term $E_t[B_{t+1}^2 U_{t+1}^{-1}] - E_t[B_{t+1}]^2$ is strictly positive using Cauchy-Schwarz. To see this, note that I can multiply the first term by $E_t\left[U_{t+1}\right] = E_t[\frac{u'(c_{t+1})}{u'(c_t)}] = 1$ without changing the expression:

$$\begin{aligned}
E_t\left[B_{t+1}^2 U_{t+1}^{-1}\right] - E_t\left[B_{t+1}\right]^2 &= E_t\left[(B_{t+1}\sqrt{U_{t+1}^{-1}})^2\right]E_t\left[(\sqrt{U_{t+1}})^2\right] - E_t\left[B_{t+1}\right]^2 \qquad \text{(A.14)} \\
&= E[x^2]E[y^2] - E[xy]^2 > 0,
\end{aligned}$$

with $x = B_{t+1}\sqrt{U_{t+1}^{-1}}$ and $y = \sqrt{U_{t+1}}$. Without the HARA utility, I could not apply Cauchy-Schwarz to sign this expression, because I would note have the same expression B in the two terms.

Time-varying demographics, discount factors, and interest rates. As before, in a model with time-varying demographic characteristics, discount factors and interest rates, the Euler equation becomes:

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

with $R_{t,t+s} \equiv \Pi_{k=1}^{s}\beta_{t+k}(1+r_{t+k-1})e^{\delta(z_{t+k}-z_{t+k-1})}$ not necessarily equal to one. Because I do not need to compare the expressions with their perfect foresight values, this does not affect the reasoning, which goes through adding an extra $R_{t,t+1}$ term along $u'(c_{t+1})$.

A.4 Proof of Lemma (iii)

Proof of Lemma (iii) with isoelastic utility. I first establish that, when $u(.)$ is isoelastic, $c_t = c_t(a_{t-1}, e^{p_t}, e^{\varepsilon_t})$ is homogeneous of degree one in a_{t-1} and e^{p_t}. The isoelastic utility case is a subcase of the set of utility functions such that relative prudence is strictly larger than one and increasing, because denoting $\rho > 0$ the relative risk aversion, relative prudence is $\rho + 1$, thus strictly larger than one and constant. By Euler's homogeneous function theorem, showing homogeneity of degree one in a_{t-1} and e^{p_t} is equivalent to showing that:

$$c_t = a_{t-1}\frac{\partial c_t}{\partial a_{t-1}} + e^{p_t}\frac{\partial c_t}{\partial e^{p_t}}.$$

I prove it by backward induction. At the last period $t = T$:

$$c_T = (1+r)a_{T-1} + e^{\varepsilon_T}e^{p_T} = \frac{\partial c_T}{\partial a_{T-1}}a_{T-1} + \frac{\partial c_T}{\partial e^{p_T}}e^{p_T}. \tag{A.15}$$

This means that $c_T(a_{T-1}, e^{p_T}, e^{\varepsilon_T})$ is homogeneous of degree one in a_{T-1} and e^{p_T}. Now, I assume that $c_{t+1}(a_t, e^{p_{t+1}}, e^{\varepsilon_{t+1}})$ is homogeneous of degree one in a_t and $e^{p_{t+1}}$ and prove that $c_t(a_{t-1}, e^{p_t}, e^{\varepsilon_t})$ must then be homogeneous of degree one in a_{t-1} and e^{p_t}. To show this, I differentiate both sides of the Euler equation (2.8) with respect to e^{p_t}:

$$\frac{\partial c_t}{\partial e^{p_t}} = E_t\big[\big(\frac{\partial a_t}{\partial e^{p_t}}\frac{\partial c_{t+1}}{\partial a_t} + \frac{\partial e^{p_{t+1}}}{\partial e^{p_t}}\frac{\partial c_{t+1}}{\partial e^{p_{t+1}}}\big)\frac{-u''(c_{t+1})}{-u''(c_t)}\big] \tag{A.16}$$

$$= E_t\big[\big((-\frac{\partial c_t}{\partial e^{p_t}} + e^{\varepsilon_t})\frac{\partial c_{t+1}}{\partial a_t} + \frac{e^{p_{t+1}}}{e^{p_t}}\frac{\partial c_{t+1}}{\partial e^{p_{t+1}}}\big)\frac{-u''(c_{t+1})}{-u''(c_t)}\big] \tag{A.17}$$

$$= \frac{E_t\big[\big(\frac{a_t-(1+r)a_{t-1}+c_t}{e^{p_t}}\frac{\partial c_{t+1}}{\partial a_t} + \frac{e^{p_{t+1}}}{e^{p_t}}\frac{\partial c_{t+1}}{\partial e^{p_{t+1}}}\big)\frac{-u''(c_{t+1})}{-u''(c_t)}\big]}{1+E_t\big[\frac{\partial c_{t+1}}{\partial a_{t+1}}\frac{-u''(c_{t+1})}{-u''(c_t)}\big]} \tag{A.18}$$

$$= \frac{\frac{1}{e^{p_t}}E_t\big[\big(a_t\frac{\partial c_{t+1}}{\partial a_t} + e^{p_{t+1}}\frac{\partial c_{t+1}}{\partial e^{p_{t+1}}}\big)\frac{-u''(c_{t+1})}{-u''(c_t)}\big] + \frac{-(1+r)a_{t-1}+c_t}{e^{p_t}}E_t\big[\frac{\partial c_{t+1}}{\partial a_t}\frac{-u''(c_{t+1})}{-u''(c_t)}\big]}{1+E_t\big[\frac{\partial c_{t+1}}{\partial a_{t+1}}\frac{-u''(c_{t+1})}{-u''(c_t)}\big]} \tag{A.19}$$

$$= \frac{\frac{c_t}{e^{p_t}}E_t\big[\frac{c_{t+1}}{c_t}\frac{-u''(c_{t+1})}{-u''(c_t)}\big] + \frac{-(1+r)a_{t-1}+c_t}{e^{p_t}}E_t\big[\frac{\partial c_{t+1}}{\partial a_t}\frac{-u''(c_{t+1})}{-u''(c_t)}\big]}{1+E_t\big[\frac{\partial c_{t+1}}{\partial a_{t+1}}\frac{-u''(c_{t+1})}{-u''(c_t)}\big]} \tag{A.20}$$

$$= \frac{\frac{c_t}{e^{p_t}} + \frac{-(1+r)a_{t-1}+c_t}{e^{p_t}}E_t\big[\frac{\partial c_{t+1}}{\partial a_t}\frac{-u''(c_{t+1})}{-u''(c_t)}\big]}{1+E_t\big[\frac{\partial c_{t+1}}{\partial a_{t+1}}\frac{-u''(c_{t+1})}{-u''(c_t)}\big]} \tag{A.21}$$

$$= \frac{c_t}{e^{p_t}} - \frac{a_{t-1}}{e^{p_t}}\frac{\partial c_t}{\partial a_{t-1}}. \tag{A.22}$$

This means that $c_t = e^{p_t}\frac{\partial c_t}{\partial e^{p_t}} + \frac{a_{t-1}}{e^{p_t}}\frac{\partial c_t}{\partial a_{t-1}}$, so $c_t(a_{t-1}, e^{p_t}, e^{\varepsilon_t})$ is homogeneous of degree one in a_{t-1} and e^{p_t}. I move from the first line to the second line by substituting $\frac{\partial a_t}{\partial e^{p_t}}$ with $-\frac{\partial c_t}{\partial e^{p_t}} + e^{\varepsilon_t}$, and substituting $\frac{\partial e^{p_{t+1}}}{\partial e^{p_t}}$ with $e^{\eta_{t+1}} = \frac{e^{p_{t+1}}}{e^{p_t}}$. I move from the second to the third line by taking the term on the right-hand side proportional to $-\frac{\partial c_t}{\partial e^{p_t}}$ to the left-hand side and rearranging, and by using the budget constraint at t to substitute e^{ε_t} with $\frac{a_t-(1+r)a_{t-1}+c_t}{e^{p_t}}$. I move from the third to the fourth line by rearranging the terms to have $a_t\frac{\partial c_{t+1}}{\partial a_t} + e^{p_{t+1}}\frac{\partial c_{t+1}}{\partial e^{p_{t+1}}}$ appear. I move from the fourth to the fifth line using that, because $c_{t+1}(a_t, e^{p_{t+1}}, e^{\varepsilon_{t+1}})$ is homogeneous of degree one then $a_t\frac{\partial c_{t+1}}{\partial a_t} + e^{p_{t+1}}\frac{\partial c_{t+1}}{\partial e^{p_{t+1}}} = c_{t+1}$, and rewriting $c_{t+1} = c_t(c_{t+1}/c_t)$. I move from the fourth to the fifth line using that, when $u'(.)$ is isoelastic:

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

I move from the fifth to the sixth line, substituting $(1+r)\frac{E_t\big[\frac{\partial c_{t+1}}{\partial a_t}\frac{-u''(c_{t+1})}{-u''(c_t)}\big]}{1+E_t\big[\frac{\partial c_{t+1}}{\partial a_{t+1}}\frac{-u''(c_{t+1})}{-u''(c_t)}\big]}$ with $\frac{\partial c_t}{\partial a_{t-1}}$.

Extension of Lemma (iii) with isoelastic utility. To prove Lemma (iv), I rely on an extension of Lemma (iii) stating that $c_t = c_t(a_{t-2}, e^{p_{t-1}}, e^{\eta_t}, e^{\varepsilon_t})$ is homogeneous of degree one in past wealth a_{t-2} and past persistent earnings $e^{p_{t-1}}$:

$$c_t = a_{t-2}\frac{\partial c_t}{\partial a_{t-2}} + e^{p_{t-1}}\frac{\partial c_t}{\partial e^{p_{t-1}}}.$$

I prove it here by backward induction. At the last period $t = T$:

$$c_T = (1+r)^2 a_{T-2} + (1+r)e^{p_{T-1}}e^{\varepsilon_{T-1}} + e^{p_{T-1}}e^{\eta_T}e^{\varepsilon_T} - (1+r)c_{T-1},$$

which implies that $\frac{\partial c_T}{\partial a_{T-2}} = (1+r)^2 - (1+r)\frac{\partial c_{T-1}}{\partial a_{T-2}}$ and $\frac{\partial c_T}{\partial e^{p_{T-1}}} = (1+r)e^{\varepsilon_{T-1}} + e^{\eta_T}e^{\varepsilon_T} - (1+r)\frac{\partial c_{T-1}}{\partial e^{p_{T-1}}}$. As a result:

$$\begin{aligned} c_T &= (1+r)^2 a_{T-2} + (1+r)e^{p_{T-1}}e^{\varepsilon_{T-1}} + e^{p_{T-1}}e^{\eta_T}e^{\varepsilon_T} - (1+r)c_{T-1} \\ &= (1+r)^2 a_{T-2} + (1+r)e^{p_{T-1}}e^{\varepsilon_{T-1}} + e^{p_{T-1}}e^{\eta_T}e^{\varepsilon_T} - (1+r)\big(a_{T-2}\frac{\partial c_{T-1}}{\partial a_{T-2}} + e^{p_{T-1}}\frac{\partial c_{T-1}}{\partial e^{p_{T-1}}}\big) \\ &= a_{T-2}\frac{\partial c_T}{\partial a_{T-2}} + e^{p_{T-1}}\frac{\partial c_T}{\partial e^{p_{T-1}}}. \end{aligned}$$

I move from the first to the second line using that, from Lemma (iii), $c_{T-1} = a_{T-2}\frac{\partial c_{T-1}}{\partial a_{T-2}} + e^{p_{T-1}}\frac{\partial c_{T-1}}{\partial e^{p_{T-1}}}$. I move from the second to the third line substituting with the expressions of $\frac{\partial c_T}{\partial a_{T-2}}$ and $\frac{\partial c_T}{\partial e^{p_{T-1}}}$.

For the same reason as in the proof of Lemma (iii), assuming this proposition is true at $t+1$, it then is true at t. To see this, one can derive both sides of the Euler equation with respect to $e^{p_{t-1}}$, and see that the partial effect of $e^{p_{t-1}}$ on $c_t = c_t(a_{t-2}, e^{p_{t-1}}, e^{\eta_t}, e^{\varepsilon_t})$ equals $\frac{c_t}{e^{p_{t-1}}} - \frac{a_{t-2}}{e^{p_{t-1}}}\frac{\partial c_t}{\partial a_{t-2}}$.

Proof of Lemma (iii) with strictly increasing relative prudence. I now show that, when relative prudence is strictly larger than one and strictly increasing, then c_t is strictly larger than the weighted sum of its derivative with respect to a_{t-1} and e^{p_t}. This means showing that at any period $t < T$:

$$c_t > a_{t-1}\frac{\partial c_t}{\partial a_{t-1}} + e^{p_t}\frac{\partial c_t}{\partial e^{p_t}}.$$

I prove it by backward induction. At the last period $t = T$, $c_T(a_{T-1}, e^{p_T}, e^{\varepsilon_T})$ is still homogeneous of degree one in wealth and in persistent earnings, because it is so regardless of the utility function. Then, I assume that $c_{t+1} \geq a_t\frac{\partial c_{t+1}}{\partial a_t} + e^{p_{t+1}}\frac{\partial c_{t+1}}{\partial e^{p_{t+1}}}$, and use it to rearrange expression

(A.19):

$$\frac{\partial c_t}{\partial e^{p_t}} = \frac{\frac{1}{e^{p_t}} E_t\left[\left(a_t \frac{\partial c_{t+1}}{\partial a_t} + e^{p_{t+1}} \frac{\partial c_{t+1}}{\partial e^{p_{t+1}}}\right) \frac{-u''(c_{t+1})}{-u''(c_t)}\right] + \frac{-(1+r)a_{t-1}+c_t}{e^{p_t}} E_t\left[\frac{\partial c_{t+1}}{\partial a_t} \frac{-u''(c_{t+1})}{-u''(c_t)}\right]}{1 + E_t\left[\frac{\partial c_{t+1}}{\partial a_{t+1}} \frac{-u''(c_{t+1})}{-u''(c_t)}\right]}$$

$$\leq \frac{\frac{c_t}{e^{p_t}} E_t\left[\frac{c_{t+1}}{c_t} \frac{-u''(c_{t+1})}{-u''(c_t)}\right] + \frac{-(1+r)a_{t-1}+c_t}{e^{p_t}} E_t\left[\frac{\partial c_{t+1}}{\partial a_t} \frac{-u''(c_{t+1})}{-u''(c_t)}\right]}{1 + E_t\left[\frac{\partial c_{t+1}}{\partial a_{t+1}} \frac{-u''(c_{t+1})}{-u''(c_t)}\right]} \tag{A.23}$$

$$< \frac{\frac{c_t}{e^{p_t}} + \frac{-(1+r)a_{t-1}+c_t}{e^{p_t}} E_t\left[\frac{\partial c_{t+1}}{\partial a_t} \frac{-u''(c_{t+1})}{-u''(c_t)}\right]}{1 + E_t\left[\frac{\partial c_{t+1}}{\partial a_{t+1}} \frac{-u''(c_{t+1})}{-u''(c_t)}\right]} \tag{A.24}$$

$$< \frac{c_t}{e^{p_t}} - \frac{a_{t-1}}{e^{p_t}} \frac{\partial c_t}{\partial a_{t-1}}. \tag{A.25}$$

I move from the first to the second line using that a non-strict version of Lemma (iii) is true at $t+1$. I move from the second to the third line using that, when relative prudence $u'''(c)c/(-u''(c))$ is strictly larger than one and strictly increasing, then $h(c) = -u''(c) \times c$ is strictly decreasing in c and the premium φ^h associated with $h(c)$ is strictly smaller than the premium φ associated with $u'(c)$. Indeed, first, $h'(c) = -u'''(c)c + (-u''(c))$ is strictly negative when $u'''(c)c/(-u''(c)) > 1$. Second, to compare the two premia, note that the premium associated with a function $u_1(.)$ is strictly smaller than the premium associated with $u_2(.)$ if $u_1'(.)/u_2'(.)$ is strictly positive and strictly increasing. Here $h'(.)/u''(.) = u'''(c)c/(-u''(c)) - 1$ is strictly positive and strictly increasing when relative prudence $u'''(c)c/(-u''(c))$ is strictly larger than one and strictly increasing. As a result, because $h(c)$ is strictly decreasing and $\varphi^h < \varphi$:

$$E_t[-u''(c_{t+1})c_{t+1}] = E_t[h(c_{t+1})] = h(E_t[c_{t+1}] - \varphi_t^h) < h(E_t[c_{t+1}] - \varphi_t) = h(c_t) = -u''(c_t)c_t.$$

This means that $E_t[-u''(c_{t+1})c_{t+1}]/(-u''(c_t)c_t) < 1$, which I use to substitute in the third line. To move from the third to the fourth line, I substitute $(1+r)\frac{E_t\left[\frac{\partial c_{t+1}}{\partial a_t} \frac{-u''(c_{t+1})}{-u''(c_t)}\right]}{1+E_t\left[\frac{\partial c_{t+1}}{\partial a_{t+1}} \frac{-u''(c_{t+1})}{-u''(c_t)}\right]}$ with $\frac{\partial c_t}{\partial a_{t-1}}$.

Similar to the homogeneity case, consumption is also strictly larger than a weighted sum of its derivative with respect to a_{t-2} and $e^{p_{t-1}}$.

Proof of Lemma (iii) with strictly decreasing relative prudence. The proof is the converse of the proof with strictly increasing relative prudence. In this case, $h(c) = -u''(c) \times c$ is still strictly decreasing in c but such that $\varphi^h > \varphi$, so:

$$E_t[-u''(c_{t+1})c_{t+1}] = E_t[h(c_{t+1})] = h(E_t[c_{t+1}] - \varphi_t^h) > h(E_t[c_{t+1}] - \varphi_t) = h(c_t) = -u''(c_t)c_t.$$

Time-varying demographics, discount factors, and interest rates. As before, in a model with time-varying demographic characteristics, discount factors and interest rates, the Euler equation becomes:

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

with $R_{t,t+s} \equiv \Pi_{k=1}^{s} \beta_{t+k}(1+r_{t+k-1})e^{\delta(z_{t+k}-z_{t+k-1})}$ not necessarily equal to one. Because I do not need to compare the expressions with their perfect foresight values, this does not affect the reasoning, which goes through adding an extra $R_{t,t+1}$ term along $u'(c_{t+1})$.

A.5 Proof of Lemma (iv)

Proof of Lemma (iv) with isoelastic utility. I first establish that, when $u(.)$ is isoelastic, the MPC function $(\partial c_t/\partial a_{t-1}) = (\partial c_t/\partial a_{t-1})(a_{t-1}, e^{p_t}, e^{\varepsilon_t})$ is homogeneous of degree zero in a_{t-1} and e^{p_t}. Again, the isoelastic utility case is a subcase of the set of utility functions such that relative prudence is strictly larger than one and increasing, because denoting $\rho > 0$ the relative risk aversion, relative prudence is $\rho + 1$, thus strictly larger than one and constant. By Euler's homogeneous function theorem, showing homogeneity of degree zero in a_{t-1} and e^{p_t} is equivalent to showing that:

$$0 = e^{p_t}\frac{\partial^2 c_t}{\partial a_{t-1} e^{p_t}} + a_{t-1}\frac{\partial^2 c_t}{\partial a_{t-1}^2}.$$

Rearranging, this also means showing that at any $t < T$:

$$\frac{\partial^2 c_t}{\partial a_{t-1} e^{p_t}} = -\frac{a_{t-1}}{e^{p_t}}\frac{\partial^2 c_t}{\partial a_{t-1}^2}.$$

A quick way to prove this is to use the fact that, from the proof of Lemma (iii) above, when utility is isoelastic:

$$c_t = a_{t-1}\frac{\partial c_t}{\partial a_{t-1}} + e^{p_t}\frac{\partial c_t}{\partial e^{p_t}}.$$

Differentiating both sides with respect to a change in a_{t-1} yields:

$$\begin{aligned}\frac{\partial c_t}{\partial a_{t-1}} &= \frac{\partial c_t}{\partial a_{t-1}} + a_{t-1}\frac{\partial^2 c_t}{\partial a_{t-1}^2} + e^{p_t}\frac{\partial^2 c_t}{\partial a_{t-1} e^{p_t}} \\ 0 &= a_{t-1}\frac{\partial^2 c_t}{\partial a_{t-1}^2} + e^{p_t}\frac{\partial^2 c_t}{\partial a_{t-1} e^{p_t}}.\end{aligned}$$

Now, this proof is unfortunately not suitable for inequalities extensions: the fact that $c_t \geq a_{t-1}\frac{\partial c_t}{\partial a_{t-1}} + e^{p_t}\frac{\partial c_t}{\partial e^{p_t}}$ does not imply by differentiation that $0 \geq a_{t-1}\frac{\partial^2 c_t}{\partial a_{t-1}^2} + e^{p_t}\frac{\partial^2 c_t}{\partial a_{t-1} e^{p_t}}$. Thus, I also take the long route to proving the result, in order to show how it works in this simpler setting and then explain more easily the inequality case. I prove it by backward induction.

At $t=T$, $\frac{\partial^2 c_t}{\partial a_{t-1}^2}=\frac{\partial^2 c_t}{\partial a_{t-1}e^{p_t}}=0$, so Lemma (iv) is true at this last period. I then assume that Lemma (iv) is true at $t+1$ to show it must then be true at t. I derive the expression of $\frac{\partial^2 c_t}{\partial a_{t-1}^2}$ by differentiating twice both sides of the Euler equation with respect to a_{t-1}, and rearrange (noting that a_{t-1} only affects c_{t+1} through its effect on a_t, and substituting $(\partial c_t/\partial a_{t-1})$ with $E_t\big[(\partial c_{t+1}/\partial a_{t-1})(-u''(c_{t+1})/-u''(c_t))\big]$):

$$\frac{\partial^2 c_t}{\partial a_{t-1}^2}(-u''(c_t))-(\frac{\partial c_t}{\partial a_{t-1}})^2u'''(c_t)$$
$$=E_t\big[\frac{\partial^2 c_{t+1}}{\partial a_{t-1}^2}(-u''(c_{t+1}))\big]-E_t\big[(\frac{\partial c_{t+1}}{\partial a_{t-1}})^2u'''(c_{t+1})\big]$$
$$\frac{\partial^2 c_t}{\partial a_{t-1}^2}(-u''(c_t))-(\frac{\partial c_t}{\partial a_{t-1}})^2u'''(c_t)$$
$$=E_t[(\frac{\partial^2 a_t}{\partial a_{t-1}^2}\frac{\partial c_{t+1}}{\partial a_t}+(\frac{\partial a_t}{\partial a_{t-1}})^2\frac{\partial^2 c_{t+1}}{\partial a_t^2})(-u''(c_{t+1}))]-E_t[(\frac{\partial c_{t+1}}{\partial a_{t-1}})^2u'''(c_{t+1})]$$
$$\frac{\partial^2 c_t}{\partial a_{t-1}^2}\left(1+E_t\left[(\frac{\partial c_{t+1}}{\partial a_t}\frac{-u''(c_{t+1})}{-u''(c_t)}\right]\right)$$
$$=E_t\big[(\frac{\partial a_t}{\partial a_{t-1}})^2\frac{\partial^2 c_{t+1}}{\partial a_t^2}\frac{-u''(c_{t+1})}{-u''(c_t)}\big]$$
$$-\frac{u'''(c_t)}{-u''(c_t)}\left(E_t\big[(\frac{\partial c_{t+1}}{\partial a_{t-1}})^2\frac{u'''(c_{t+1})}{u'''(c_t)}\big]-E_t\left[\frac{\partial c_{t+1}}{\partial a_{t-1}}\frac{-u''(c_{t+1})}{-u''(c_t)}\right]^2\right)$$

I now differentiate both sides of the Euler equation twice, first with respect to a_{t-1}, second with respect to e^{p_t}:

$$\frac{\partial^2 c_t}{\partial a_{t-1}e^{p_t}}(-u''(c_t))-(\frac{\partial c_t}{\partial a_{t-1}})(\frac{\partial c_t}{\partial e^{p_t}})u'''(c_t)$$
$$=E_t\big[\frac{\partial^2 c_{t+1}}{\partial a_{t-1}e^{p_t}}(-u''(c_{t+1}))\big]-E_t\big[(\frac{\partial c_{t+1}}{\partial a_{t-1}})(\frac{\partial c_{t+1}}{\partial e^{p_t}})u'''(c_{t+1})\big]$$
$$\frac{\partial^2 c_t}{\partial a_{t-1}e^{p_t}}(-u''(c_t))-(\frac{\partial c_t}{\partial a_{t-1}})(\frac{\partial c_t}{\partial e^{p_t}})u'''(c_t)$$
$$=E_t[\left(\frac{\partial^2 a_t}{\partial a_{t-1}e^{p_t}}\frac{\partial c_{t+1}}{\partial a_t}+\frac{\partial a_t}{\partial a_{t-1}}(\frac{\partial a_t}{e^{p_t}}\frac{\partial^2 c_{t+1}}{\partial a_t^2}+\frac{\partial e^{p_{t+1}}}{e^{p_t}}\frac{\partial^2 c_{t+1}}{\partial a_t e^{p_{t+1}}})\right)(-u''(c_{t+1}))]$$
$$-E_t[(\frac{\partial c_{t+1}}{\partial a_{t-1}})^2u'''(c_{t+1})]$$
$$\frac{\partial^2 c_t}{\partial a_{t-1}e^{p_t}}\left(1+E_t\left[(\frac{\partial c_{t+1}}{\partial a_t}\frac{-u''(c_{t+1})}{-u''(c_t)}\right]\right)$$
$$=E_t\big[\frac{\partial a_t}{\partial a_{t-1}}(\frac{\partial a_t}{\partial e^{p_t}}\frac{\partial^2 c_{t+1}}{\partial a_t^2}+\frac{\partial e^{p_{t+1}}}{\partial e^{p_t}}\frac{\partial^2 c_{t+1}}{\partial a_t e^{p_{t+1}}})\frac{-u''(c_{t+1})}{-u''(c_t)}\big]$$
$$-\frac{u'''(c_t)}{-u''(c_t)}\left(E_t\big[\frac{\partial c_{t+1}}{\partial a_{t-1}}\frac{\partial c_{t+1}}{\partial e^{p_t}}\frac{u'''(c_{t+1})}{u'''(c_t)}\big]-E_t\left[\frac{\partial c_{t+1}}{\partial a_{t-1}}\frac{-u''(c_{t+1})}{-u''(c_t)}\right]E_t\left[\frac{\partial c_{t+1}}{\partial e^{p_t}}\frac{-u''(c_{t+1})}{-u''(c_t)}\right]\right)$$

I then proceed in two steps. First, I show that, when Lemma (iv) is true at $t+1$, then the term in red in the expression of $\frac{\partial^2 c_t}{\partial a_{t-1} e^{p_t}}$, which I denote A_t^{ap}, is equal to $-(a_{t-1}/e^{p_t})$ times the term in pink in the expression of $\frac{\partial^2 c_t}{\partial a_{t-1}^2}$, which I denote A_t^{aa}. Second, I show that, when Lemma (iv) is true at $t+1$, then the term in blue in the expression of $\frac{\partial^2 c_t}{\partial a_{t-1} e^{p_t}}$, which I denote B_t^{ap}, is equal to $-(a_{t-1}/e^{p_t})$ times the term in light blue in the expression of $\frac{\partial^2 c_t}{\partial a_{t-1}^2}$, which I denote B_t^{aa}. In A_t^{ap}, I substitute $\partial^2 c_{t+1}/\partial a_t e^{p_{t+1}}$ with $-(a_t/e^{p_{t+1}})(\partial^2 c_{t+1}/\partial a_t^2)$:

$$
\begin{aligned}
A_t^{ap} &= E_t\Big[\frac{\partial a_t}{\partial a_{t-1}}\Big(\frac{\partial a_t}{\partial e^{p_t}}\frac{\partial^2 c_{t+1}}{\partial a_t^2} + \frac{\partial e^{p_{t+1}}}{\partial e^{p_t}}\frac{\partial^2 c_{t+1}}{\partial a_t e^{p_{t+1}}}\Big)\frac{-u''(c_{t+1})}{-u''(c_t)}\Big] \\
&= E_t\Big[\frac{\partial a_t}{\partial a_{t-1}}\Big(\frac{\partial a_t}{\partial e^{p_t}} - \frac{\partial e^{p_{t+1}}}{\partial e^{p_t}}\frac{a_t}{e^{p_{t+1}}}\Big)\frac{\partial^2 c_{t+1}}{\partial a_t^2}\frac{-u''(c_{t+1})}{-u''(c_t)}\Big] \\
&= E_t\Big[\frac{\partial a_t}{\partial a_{t-1}}\Big(e^{\varepsilon_t} - \frac{\partial c_t}{\partial e^{p_t}} - \frac{e^{p_{t+1}}}{e^{p_t}}\frac{a_t}{e^{p_{t+1}}}\Big)\frac{\partial^2 c_{t+1}}{\partial a_t^2}\frac{-u''(c_{t+1})}{-u''(c_t)}\Big] \\
&= E_t\Big[\frac{\partial a_t}{\partial a_{t-1}}\Big(e^{\varepsilon_t} - \Big(\frac{c_t}{e^{p_t}} - \frac{a_{t-1}}{e^{p_t}}\frac{\partial c_t}{\partial a_{t-1}}\Big) - \frac{e^{p_{t+1}}}{e^{p_t}}\frac{a_t}{e^{p_{t+1}}}\Big)\frac{\partial^2 c_{t+1}}{\partial a_t^2}\frac{-u''(c_{t+1})}{-u''(c_t)}\Big] \\
&= E_t\Big[\frac{\partial a_t}{\partial a_{t-1}}\Big(\frac{a_t-(1+r)a_{t-1}+c_t}{e^{p_t}} - \Big(\frac{c_t}{e^{p_t}} - \frac{a_{t-1}}{e^{p_t}}\frac{\partial c_t}{\partial a_{t-1}}\Big) - \frac{a_t}{e^{p_t}}\Big)\frac{\partial^2 c_{t+1}}{\partial a_t^2}\frac{-u''(c_{t+1})}{-u''(c_t)}\Big] \\
&= -\frac{a_{t-1}}{e^{p_t}}E_t\Big[\frac{\partial a_t}{\partial a_{t-1}}\Big((1+r) - \frac{\partial c_t}{\partial a_{t-1}}\Big)\frac{\partial^2 c_{t+1}}{\partial a_t^2}\frac{-u''(c_{t+1})}{-u''(c_t)}\Big] \\
&= -\frac{a_{t-1}}{e^{p_t}}E_t\Big[\Big(\frac{\partial a_t}{\partial a_{t-1}}\Big)^2\frac{\partial^2 c_{t+1}}{\partial a_t^2}\frac{-u''(c_{t+1})}{-u''(c_t)}\Big] = -\frac{a_{t-1}}{e^{p_t}}A_t^{aa}.
\end{aligned}
$$

Since the extension of Lemma (iii) implies that $c_t = c_t(a_{t-2}, e^{p_{t-1}}, e^{\eta_t}, e^{\varepsilon_t})$ is homogeneous of degree one in a_{t-2} and $e^{p_{t-1}}$ when utility is isoelastic, in B_t^{ap}, I substitute $(\partial c_{t+1}/\partial e^{p_t})$ with $(c_{t+1}/e^{p_t}) - (a_{t-1}/e^{p_t})(\partial c_{t+1}/\partial a_{t-1})$:

$$
\begin{aligned}
B_t^{ap} &= -\frac{u'''(c_t)}{-u''(c_t)}\left(E_t\Big[\frac{\partial c_{t+1}}{\partial a_{t-1}}\frac{\partial c_{t+1}}{\partial e^{p_t}}\frac{u'''(c_{t+1})}{u'''(c_t)}\Big] - E_t\left[\frac{\partial c_{t+1}}{\partial a_{t-1}}\frac{-u''(c_{t+1})}{-u''(c_t)}\right]E_t\left[\frac{\partial c_{t+1}}{\partial e^{p_t}}\frac{-u''(c_{t+1})}{-u''(c_t)}\right]\right) \\
&= \frac{a_{t-1}}{e^{p_t}}\frac{u'''(c_t)}{-u''(c_t)}\left(E_t\Big[\Big(\frac{\partial c_{t+1}}{\partial a_{t-1}}\Big)^2\frac{u'''(c_{t+1})}{u'''(c_t)}\Big] - E_t\left[\frac{\partial c_{t+1}}{\partial a_{t-1}}\frac{-u''(c_{t+1})}{-u''(c_t)}\right]^2\right) \\
&- \frac{u'''(c_t)}{-u''(c_t)}\left(E_t\Big[\frac{\partial c_{t+1}}{\partial a_{t-1}}\frac{c_{t+1}}{e^{p_t}}\frac{u'''(c_{t+1})}{u'''(c_t)}\Big] - E_t\left[\frac{\partial c_{t+1}}{\partial a_{t-1}}\frac{-u''(c_{t+1})}{-u''(c_t)}\right]E_t\left[\frac{c_{t+1}}{e^{p_t}}\frac{-u''(c_{t+1})}{-u''(c_t)}\right]\right) \\
&= -\frac{a_{t-1}}{e^{p_t}}B_t^{aa} - \frac{u'''(c_t)}{-u''(c_t)}\frac{1}{e^{p_t}}\left(E_t\Big[\frac{\partial c_{t+1}}{\partial a_{t-1}}\frac{c_{t+1}}{c_t}\frac{u'''(c_{t+1})}{u'''(c_t)}\Big] - E_t\left[\frac{\partial c_{t+1}}{\partial a_{t-1}}\frac{-u''(c_{t+1})}{-u''(c_t)}\right]E_t\left[\frac{c_{t+1}}{c_t}\frac{-u''(c_{t+1})}{-u''(c_t)}\right]\right) \\
&= -\frac{a_{t-1}}{e^{p_t}}B_t^{aa} - \frac{u'''(c_t)}{-u''(c_t)}\frac{1}{e^{p_t}}\left(E_t\Big[\frac{\partial c_{t+1}}{\partial a_{t-1}}\frac{-u''(c_{t+1})}{-u''(c_t)}\Big] - E_t\left[\frac{\partial c_{t+1}}{\partial a_{t-1}}\frac{-u''(c_{t+1})}{-u''(c_t)}\right]E_t\left[\frac{u'(c_{t+1})}{u'(c_t)}\right]\right) \\
&= -\frac{a_{t-1}}{e^{p_t}}B_t^{aa}.
\end{aligned}
$$

As a result, it must be that:

$$\frac{\partial^2 c_t}{\partial a_{t-1} e^{p_t}} = -\frac{a_{t-1}}{e^{p_t}} \frac{\partial^2 c_t}{\partial a_{t-1}^2}.$$

Proof of Lemma (iv) with strictly increasing relative prudence. I now consider the case with a utility function displaying HARA, and a relative prudence strictly larger than one and strictly increasing. I show that, in that case, the weighted sum of the partial effects of wealth and persistent earnings on the MPC is strictly larger than zero. This means showing that at any $t < T$:

$$0 > e^{p_t} \frac{\partial^2 c_t}{\partial a_{t-1} e^{p_t}} + a_{t-1} \frac{\partial^2 c_t}{\partial a_{t-1}^2}.$$

Rearranging, this also means showing that at any $t < T$:

$$\frac{\partial^2 c_t}{\partial a_{t-1} e^{p_t}} < -\frac{a_{t-1}}{e^{p_t}} \frac{\partial^2 c_t}{\partial a_{t-1}^2}.$$

A non-strict version of Lemma (iv) is true at $t = T$ because $\frac{\partial^2 c_T}{\partial a_{T-1}^2} = \frac{\partial^2 c_T}{\partial a_{T-1} e^{p_T}} = 0$. Now I assume that a non-strict version of Lemma (iv) is true at at $t+1$ and prove that it must then be true at t. Using Lemma (iv) at $t+1$, the term A_t^{ap} becomes:

$$\begin{aligned}
A_t^{ap} &= E_t\Big[\frac{\partial a_t}{\partial a_{t-1}}\Big(\frac{\partial a_t}{\partial e^{p_t}} \frac{\partial^2 c_{t+1}}{\partial a_t^2} + \frac{\partial e^{p_{t+1}}}{\partial e^{p_t}} \frac{\partial^2 c_{t+1}}{\partial a_t e^{p_{t+1}}}\Big) \frac{-u''(c_{t+1})}{-u''(c_t)}\Big] \\
&\leq E_t\Big[\frac{\partial a_t}{\partial a_{t-1}}\Big(\frac{\partial a_t}{\partial e^{p_t}} - \frac{\partial e^{p_{t+1}}}{\partial e^{p_t}} \frac{a_t}{e^{p_{t+1}}}\Big) \frac{\partial^2 c_{t+1}}{\partial a_t^2} \frac{-u''(c_{t+1})}{-u''(c_t)}\Big] \\
&\leq E_t\Big[\frac{\partial a_t}{\partial a_{t-1}}\Big(e^{\varepsilon_t} - \frac{\partial c_t}{\partial e^{p_t}} - \frac{e^{p_{t+1}}}{e^{p_t}} \frac{a_t}{e^{p_{t+1}}}\Big) \frac{\partial^2 c_{t+1}}{\partial a_t^2} \frac{-u''(c_{t+1})}{-u''(c_t)}\Big] \\
&< E_t\Big[\frac{\partial a_t}{\partial a_{t-1}}\Big(e^{\varepsilon_t} - \Big(\frac{c_t}{e^{p_t}} - \frac{a_{t-1}}{e^{p_t}} \frac{\partial c_t}{\partial a_{t-1}}\Big) - \frac{e^{p_{t+1}}}{e^{p_t}} \frac{a_t}{e^{p_{t+1}}}\Big) \frac{\partial^2 c_{t+1}}{\partial a_t^2} \frac{-u''(c_{t+1})}{-u''(c_t)}\Big] \\
&< E_t\Big[\frac{\partial a_t}{\partial a_{t-1}}\Big(\frac{a_t - (1+r)a_{t-1} + c_t}{e^{p_t}} - \Big(\frac{c_t}{e^{p_t}} - \frac{a_{t-1}}{e^{p_t}} \frac{\partial c_t}{\partial a_{t-1}}\Big) - \frac{a_t}{e^{p_t}}\Big) \frac{\partial^2 c_{t+1}}{\partial a_t^2} \frac{-u''(c_{t+1})}{-u''(c_t)}\Big] \\
&< -\frac{a_{t-1}}{e^{p_t}} E_t\Big[\frac{\partial a_t}{\partial a_{t-1}}\Big((1+r) - \frac{\partial c_t}{\partial a_{t-1}}\Big) \frac{\partial^2 c_{t+1}}{\partial a_t^2} \frac{-u''(c_{t+1})}{-u''(c_t)}\Big] \\
&< -\frac{a_{t-1}}{e^{p_t}} E_t\Big[\Big(\frac{\partial a_t}{\partial a_{t-1}}\Big)^2 \frac{\partial^2 c_{t+1}}{\partial a_t^2} \frac{-u''(c_{t+1})}{-u''(c_t)}\Big] = -\frac{a_{t-1}}{e^{p_t}} A_t^{aa}.
\end{aligned}$$

The difference with the previous case is that I move the first to the second line using that, when Lemma (iv) is true at $t+1$, $\frac{\partial^2 c_{t+1}}{\partial a_t e^{p_{t+1}}} \leq -\frac{a_t}{e^{p_{t+1}}} \frac{\partial^2 c_{t+1}}{\partial a_t^2}$. I also move from the third to the fourth line using that, from Lemma (iii), when relative prudence is strictly increasing $\frac{\partial c_t}{\partial e^{p_t}} < \frac{c_t}{e^{p_t}} - \frac{a_{t-1}}{e^{p_t}} \frac{\partial c_t}{\partial a_{t-1}}$: because utility is HARA, $\frac{\partial^2 c_{t+1}}{\partial a_t^2} < 0$, and $-\frac{\partial c_t}{\partial e^{p_t}} \frac{\partial^2 c_{t+1}}{\partial a_t^2}$ is strictly smaller than

$-(\frac{c_t}{e^{p_t}} - \frac{a_{t-1}}{e^{p_t}}\frac{\partial c_t}{\partial a_{t-1}})\frac{\partial^2 c_{t+1}}{\partial a_t^2}$. I now consider the term B_t^{ap}. I defined the function $G_{t+1}(x)$:

$$G_{t+1}(x) = x \times \Big(\frac{\partial c_{t+1}}{\partial a_{t-1}}\frac{-u''(c_{t+1})}{-u''(c_t)}\frac{(u'(c_{t+1}))^{-1}}{(u'(c_t))^{-1}} - E_t\big[\frac{\partial c_{t+1}}{\partial a_{t-1}}\frac{-u''(c_{t+1})}{-u''(c_t)}\big]\Big),$$

with $\frac{\partial c_{t+1}}{\partial a_{t-1}}$, c_{t+1} and c_t the outcomes of the maximization problem described by (2.1)-(2.6). I compute the expected value at t of the first derivative of $G_{t+1}(.)$:

$$\begin{aligned}
E_t\big[G'_{t+1}(x)\big] &= E_t\big[\frac{\partial c_{t+1}}{\partial a_{t-1}}\frac{-u''(c_{t+1})}{-u''(c_t)}\frac{(u'(c_{t+1}))^{-1}}{(u'(c_t))^{-1}}\big] - E_t\big[\frac{\partial c_{t+1}}{\partial a_{t-1}}\frac{-u''(c_{t+1})}{-u''(c_t)}\big] \\
&= E_t\big[\frac{\partial c_{t+1}}{\partial a_{t-1}}\frac{-u''(c_{t+1})}{-u''(c_t)}\frac{(u'(c_{t+1}))^{-1}}{(u'(c_t))^{-1}}\big]E_t\big[\frac{u'(c_{t+1})}{u'(c_t)}\big] - E_t\big[\frac{\partial c_{t+1}}{\partial a_{t-1}}\frac{-u''(c_{t+1})}{-u''(c_t)}\big] \\
&= E_t\big[\frac{\partial c_{t+1}}{\partial a_{t-1}}\frac{-u''(c_{t+1})}{-u''(c_t)}\big] - cov_t\big(\frac{\partial c_{t+1}}{\partial a_{t-1}}\frac{-u''(c_{t+1})/u'(c_{t+1})}{-u''(c_t)/u'(c_t)}, \frac{u'(c_{t+1})}{u'(c_t)}\big) \\
&- E_t\big[\frac{\partial c_{t+1}}{\partial a_{t-1}}\frac{-u''(c_{t+1})}{-u''(c_t)}\big] \\
&< 0
\end{aligned}$$

I move from the first to the second line using that $E_t\big[\frac{u'(c_{t+1})}{u'(c_t)}\big] = 1$. I move from the second to the third line using that $E[X]E[Y] = E[XY] - cov(X,Y)$. I move from the third to the fourth line noting that, when utility is HARA, then $u'''(.)/(-u''(.)) = k(-u''(.))/u'(.)$ so decreasing relative prudence implies decreasing relative risk aversion and $(-u''(c_{t+1}))/u'(c_{t+1})$ decreases with the shocks that raise c_{t+1} (and vice-versa), while $u'(c_{t+1})$ also decreases with the shocks that raise c_{t+1} (and vice-versa). The term $\frac{\partial c_{t+1}}{\partial a_{t-1}}$ decreases with the transitory shocks that raise c_{t+1} (and vice-versa). Finally, a straightforward corollary to Theorem (ii) is that, when relative prudence is strictly above one and strictly increasing, the MPC $\frac{\partial c_{t+1}}{\partial a_{t-1}}$ decreases with the persistent shocks that raise c_{t+1} (and vice-versa). Therefore, the covariance $cov_t\big(\frac{\partial c_{t+1}}{\partial a_{t-1}}\frac{-u''(c_{t+1})/u'(c_{t+1})}{-u''(c_t)/u'(c_t)}, \frac{u'(c_{t+1})}{u'(c_t)}\big)$ must be strictly positive. The expression of $G_{t+1}(.)$ implies that :

$$B_t^{ap} = -\frac{u'''(c_t)}{-u''(c_t)}E_t[G_{t+1}(\frac{\partial c_{t+1}}{\partial e^{p_t}}\frac{-u''(c_{t+1})}{-u''(c_t)})].$$

Recall that B_t^{aa} is such that:

$$-\frac{a_{t-1}}{e^{p_t}}B_t^{aa} = -\frac{u'''(c_t)}{-u''(c_t)}E_t[G_{t+1}((\frac{c_{t+1}}{e^{p_{t+1}}} - \frac{a_t}{e^{p_{t+1}}}\frac{\partial c_{t+1}}{\partial a_t})\frac{-u''(c_{t+1})}{-u''(c_t)})].$$

Because $E_t\big[G'_{t+1}(x)\big] < 0$ and because $\frac{\partial c_{t+1}}{\partial e^{p_t}} < (\frac{c_{t+1}}{e^{p_{t+1}}} - \frac{a_t}{e^{p_{t+1}}}\frac{\partial c_{t+1}}{\partial a_t})$ when Lemma (iv) is true at

$t+1$, then:

$$B_t^{ap} = -\frac{u'''(c_t)}{-u''(c_t)} E_t[G_{t+1}(\frac{\partial c_{t+1}}{\partial e^{p_t}} \frac{-u''(c_{t+1})}{-u''(c_t)})]$$
$$< -\frac{u'''(c_t)}{-u''(c_t)} E_t[G_{t+1}((\frac{c_{t+1}}{e^{p_{t+1}}} - \frac{a_t}{e^{p_{t+1}}} \frac{\partial c_{t+1}}{\partial a_t}) \frac{-u''(c_{t+1})}{-u''(c_t)})] = -\frac{a_{t-1}}{e^{p_t}} B_t^{aa}$$

Because $A_t^{ap} < -\frac{a_{t-1}}{e^{p_t}} A_t^{aa}$ and $B_t^{ap} < -\frac{a_{t-1}}{e^{p_t}} B_t^{aa}$, it must be that:

$$\frac{\partial^2 c_t}{\partial a_{t-1} e^{p_t}} < -\frac{a_{t-1}}{e^{p_t}} \frac{\partial^2 c_t}{\partial a_{t-1}^2}.$$

Proof of Lemma (iv) with decreasing relative prudence. The proof is the converse of the proof with strictly increasing relative prudence. In that case, when the Lemma (iv) is true at $t+1$, $\frac{\partial^2 c_{t+1}}{\partial a_t e^{p_{t+1}}} \geq -\frac{a_t}{e^{p_{t+1}}} \frac{\partial^2 c_{t+1}}{\partial a_t^2}$, and from Lemma (iii), when relative prudence is strictly decreasing $\frac{\partial c_t}{\partial e^{p_t}} > \frac{c_t}{e^{p_t}} - \frac{a_{t-1}}{e^{p_t}} \frac{\partial c_t}{\partial a_{t-1}}$.

Time-varying demographics, discount factors, and interest rates. As before, in a model with time-varying demographic characteristics, discount factors and interest rates, the Euler equation becomes:

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

with $R_{t,t+s} \equiv \Pi_{k=1}^{s} \beta_{t+k}(1+r_{t+k-1}) e^{\delta(z_{t+k}-z_{t+k-1})}$ not necessarily equal to one. Because I do not need to compare the expressions with their perfect foresight values, this does not affect the reasoning, which goes through adding an extra $R_{t,t+1}$ term along $u'(c_{t+1})$.

A.6 General conditions

Model. The mechanism at play in the standard life-cycle model might hold true more generally when a change in persistent earnings shifts the precautionary motive because it modifies parameters or variables other than current and future earnings one for one. I now simply assume that, at each period t, people decide on their consumption, denoted c_t, which they finance out of $N+1$ types of revenue or resources. I assume the first corresponds to risk-free liquid wealth, denoting its level a_t^1, and $N+1$th correspond to the main component of earnings from people's main job, denoting its level e_t^p, without taking an exact stand on the shape of the earnings process. I denote a_t^2, ... a_t^N the quantities of the other resources. Although these variables are indexed by t, some of them might be determined earlier (for instance a_t^1 might be determined at $t-1$). I define the MPC as the response of consumption to a change in a_{t-1}^1. People's total resources might also depend on a vector of other components, denoted e^{ε_t}, that may affect the extent to which a_t^1, ..., a_t^N and e^{p_t} translate into a given amount of resources but would not generate resources by themselves if the others are all equal to zero. The consumption function is:

$$c_t = c_t(a_t^1, ..., a_t^N, e^{p_t}, e^{\varepsilon_t}). \tag{A.26}$$

People might be solving an intertemporal problem, in which case the uncertain future matters for their choice of c_t. I assume there exists at least one distribution of shocks, such that, under that distribution, people no longer face any uncertainty about the future, and their consumption is homogeneous of degree one in a_t^1, ..., a_t^N, e^{p_t}, and linear in a_t^1. I refer to the decisions made under this distribution as decisions under perfect foresight, and denote variables decided under perfect foresight with a superscript PF_t:

$$c_t^{PF_t} = a_t^1 \frac{\partial c_t^{PF_t}}{\partial a_t^1} + ... + a_t^T \frac{\partial c_t^{PF_t}}{\partial a_t^N} + e^{p_t} \frac{\partial c_t^{PF_t}}{\partial e^{p_t}} \quad \text{and} \quad \frac{\partial^2 c_t^{PF_t}}{\partial (a_t^1)^2} = ... = \frac{\partial^2 c_t^{PF_t}}{\partial a_t^1 a_t^N} = \frac{\partial^2 c_t^{PF_t}}{\partial a_t^1 e^{p_t}} = 0.$$

I refer to the difference between consumption under perfect foresight and consumption as precautionary saving. Indeed, since under this definition of perfect foresight people no longer face any uncertainty, I can choose it as my no-uncertainty benchmark and define precautionary saving as the distance from it.

Conditions. I show that I only require a limited set of proximate conditions for the propositions (i) and (ii) of the Theorem to hold. This set is the following:

(a) Consumption is strictly larger than it would be under perfect foresight, and there exists a non-empty set S_t^b of portfolios $(a_t^1,...a_t^N)$ such that an increase in resources that multiplies each of the $a_t^1,...,a_t^N$ by the same amount, to keep their relative size unchanged, raises consumption more than it would under perfect foresight (so it reduces precautionary saving):

$$c_t < c_t^{PF_t} \quad \text{and} \quad a_t^1 \frac{\partial c_t}{\partial a_t^1} + ... + a_t^N \frac{\partial c_t}{\partial a_t^N} > a_t^1 \frac{\partial c_t^{PF_t}}{\partial a_t^1} + ... + a_t^N \frac{\partial c_t^{PF_t}}{\partial a_t^N} \text{ when } (a_t^1,...a_t^N) \in S_t^b \; .$$

(b) There exists a non-empty set S_t^c of portfolios $(a_t^1,...a_t^N)$ such that an increase in resources that multiplies each of the $a_t^1,...,a_t^N$ by the same amount, to keep their relative size unchanged, reduces the MPC (more than it would under perfect foresight since the MPC does not vary with resources under perfect foresight):

$$a_t^1 \frac{\partial^2 c_t}{\partial (a_t^1)^2} + ... + a_t^N \frac{\partial^2 c_t}{\partial a_t^1 a_t^N} < \frac{\partial^2 c_t^{PF_t}}{\partial (a_t^1)^2} + ... + a_t^N \frac{\partial^2 c_t^{PF_t}}{\partial a_t^1 a_t^N} \text{ when } (a_t^1,...a_t^N) \in S_t^c$$

(c) The weighted sum of the partial effects of the each of the $a_t^1,...,a_t^N$ resource on consumption is smaller than consumption:

$$c_t \geq a_t^1 \frac{\partial c_t}{\partial a_t^1} + ... + a_t^T \frac{\partial c_t}{\partial a_t^N} + e^{p_t} \frac{\partial c_t}{\partial e^{p_t}}.$$

(d) The weighted sum of the partial effectsof the each of the a_t^1,...,a_t^N resource on the MPC is larger than zero:

$$0 \leq a_t^1 \frac{\partial^2 c_t}{\partial (a_t^1)^2} + ... + a_t^N \frac{\partial^2 c_t}{\partial a_t^1 a_t^N} + e^{p_t} \frac{\partial^2 c_t}{\partial a_{t-1} e^{p_t}}.$$

Generalization of Theorem (i). In the set-up described above, when conditions (a) and (c) hold, consumption increases strictly less with persistent earnings than it would under perfect foresight for consumers whose portfolios (a_t^1,...a_t^N) are in S_t^b (and precautionary saving increases with persistent earnings):

$$\frac{\partial c_t}{\partial e^{p_t}} < \frac{\partial c_t^{PF_t}}{\partial e^{p_t}} \quad \text{when } (a_t^1,...a_t^N) \in S_t^b$$

Proof of the generalization of Theorem (i). Rearranging condition (c), the partial effect of persistent earnings on consumption is:

$$\frac{\partial c_t}{\partial e^{p_t}} \leq \frac{c_t}{e^{p_t}} - \frac{a_t^1}{e^{p_t}} \frac{\partial c_t}{\partial a_t^1} - ... - \frac{a_t^N}{e^{p_t}} \frac{\partial c_t}{\partial a_t^N} < \frac{c_t^{PF_t}}{e^{p_t}} - \frac{a_{t-1}}{e^{p_t}} \frac{\partial c_t^{PF_t}}{\partial a_t^1} - ... - \frac{a_t^N}{e^{p_t}} \frac{\partial^{PF_t}}{\partial a_t^N} = \frac{\partial c_t^{PF_t}}{\partial e^{p_t}}. \quad \text{(A.27)}$$

I move from the second to the third expression using that $\frac{c_t}{e^{p_t}} < \frac{c_t^{PF_t}}{e^{p_t}}$, because people consume strictly less than they would under perfect foresight (condition (a)), and that $a_t^1 \frac{\partial c_t}{\partial a_t^1} + ... + a_t^N \frac{\partial c_t}{\partial a_t^N} > a_t^1 \frac{\partial c_t^{PF_t}}{\partial a_t^1} + ... + a_t^N \frac{\partial c_t^{PF_t}}{\partial a_t^N}$ (condition (a)). I move from the third to the fourth expression using that under perfect foresight, consumption is homogeneous of degree one a_t^1, ..., a_t^N and in e^{p_t}.

Intuitively, when consumption is homogeneous, an increase in persistent earnings is equivalent to scaling up the consumers' problem and then decreasing back all the other types of resources. The scaling up leads consumption to increase by $\frac{c_t}{e^{p_t}}$, which is smaller in the presence of uncertainty than under perfect foresight. The decrease of the other sources of revenue induces people facing uncertainty to decrease more their consumption than they would under perfect foresight (or to increase less their consumption than they would under perfect foresight if for some reason the decrease in the other sources of revenue raises consumption). As a result, their consumption responds less to an increase in persistent earnings than it would under perfect foresight and their precautionary saving increases with persistent earnings.

Generalization of Theorem (ii). In the set-up described above, when conditions (c) and (d) hold, the MPC increases strictly more with persistent earnings than it would under perfect foresight, which means that it increases strictly, for consumers whose portfolios $(a_t^1,...a_t^N)$ are in S_t^c:

$$\frac{\partial^2 c_t}{\partial a_t^1 e^{p_t}} > \frac{\partial^2 c_t^{PF_t}}{\partial a_t^1 e^{p_t}} = 0 \quad \text{when } (a_t^1,...a_t^N) \in S_t^c$$

Proof of the generalization of Theorem (ii). Rearranging condition (d), the partial effect of persistent earnings on the MPC is:

$$\frac{\partial^2 c_t}{\partial a_{t-1} e^{p_t}} \geq -\frac{a_t^1}{e^{p_t}} \frac{\partial^2 c_t}{\partial (a_t^1)^2} - ... - \frac{a_t^N}{e^{p_t}} \frac{\partial^2 c_t}{\partial a_t^1 a_t^N} > -\frac{a_t^1}{e^{p_t}} \frac{\partial^2 c_t^{PF_t}}{\partial (a_t^1)^2} - ... - \frac{a_t^N}{e^{p_t}} \frac{\partial^2 c_t^{PF_t}}{\partial a_t^1 a_t^N} = 0. \quad \text{(A.28)}$$

I move from the second to the third expression using that the effect of the other types of resources on the MPC is smaller than it would be under perfect foresight, $a_t^1 \frac{\partial^2 c_t}{\partial (a_t^1)^2} + ... + a_t^N \frac{\partial^2 c_t}{\partial a_t^1 a_t^N} < \frac{\partial^2 c_t^{PF_t}}{\partial (a_t^1)^2} + ... + a_t^N \frac{\partial^2 c_t^{PF_t}}{\partial a_t^1 a_t^N}$ (condition (a)). I move from the third to the fourth expression using that, under perfect foresight, the MPC is unaffected by people's level of resources.

Again, an increase in persistent earnings is equivalent to scaling up the consumers' problem and then shifting back the other types of resources. The scaling up does not affect the MPC, but the shift back in resources increases it more than it would under perfect foresight (because a shift in resources in the other direct raises the MPC more than it would under perfect foresight). Because the MPC does not vary with resources under perfect foresight, an increase in persistent earnings that raises the MPC more than it would under perfect foresight simply raises the MPC.

Implication. Since the variables play symmetric roles in the reasoning above, if one observes that persistent earnings raises people's MPC, then, under some conditions, this means that their persistent earnings is relatively more risky (in the sense that it strengthens precautionary behavior more) than the other resources they use to finance their consumption at this level of resources. More precisely, if condition (d) hold true, but condition (b) is replaced by the result from Theorem (ii), denoted condition (b'):

(b') The effect of persistent earnings e^{p_t} on the MPC is strictly smaller than it would be under perfect foresight:

$$\frac{\partial^2 c_t}{\partial a_t^1 e^{p_t}} < \frac{\partial^2 c_t^{PF_t}}{\partial a_t^1 e^{p_t}},$$

then condition (b) must be true, which means than increase in the $(a_t^1,...a_t^N)$ that keeps the relative shares of each resources the same reduces precautionary saving. To see this, note that:

$$\frac{a_t^1}{e^{p_t}} \frac{\partial^2 c_t}{\partial (a_t^1)^2} + ... + \frac{a_t^N}{e^{p_t}} \frac{\partial^2 c_t}{\partial a_t^1 a_t^N} \geq -\frac{\partial^2 c_t}{\partial a_{t-1} e^{p_t}} > -\frac{\partial^2 c_t^{PF_t}}{\partial a_t^1 a_t^N} = \frac{a_t^1}{e^{p_t}} \frac{\partial^2 c_t^{PF_t}}{\partial (a_t^1)^2} + ... + \frac{a_t^N}{e^{p_t}}. \quad \text{(A.29)}$$

B Data and building persistent earnings

B.1 Matching of the SCE modules

The Core survey takes place every month, but the different modules of the survey do not. They usually take place every four months or every year, at different months of the year for different modules. The three main SCE modules that I use are the Labor Market module (for current income, future expected income in four months, and the probability to be non-employed in four months), the Household Spending module (for the MPCs), and the Housing module (for the wealth categories).

The Labor Market module takes place every four months, in March, July, and November of each year. The Household Spending Module takes place every four months in April, August, and December. Finally, the Housing module takes place every year in February. As a result, I match the observations of income-related variables in November, with the observations of consumption related variables in December, and the observations of wealth-related variables in February of the next year. An additional advantage of using the survey questions that are reported at the end of the year is that the questions about the situation 'four months from now' correspond to questions about the next calendar year.

B.2 Description of the main variables

Current annual earnings. My measure of current annual earnings is the answer to the question 'How much do you make before taxes and other deductions at your [main/current] job, on an annual basis? Please include any bonuses, overtime pay, tips or commissions.' (question L4) that is in the Labor Market module of the SCE. The answer of this question is referred to as 'annual earnings' later on in the survey.[33]

I consider that the answer to this question corresponds to the respondent's current monthly earnings extrapolated backward and forward over the calendar year (say monthly earnings have a monthly component, then extrapolating would be substituting this monthly component with that of all other months and summing to get an annual measure of earnings). I can equivalently assume that the answer to the question corresponds to the respondent's earnings over the calendar year during which the question is asked. I deflate the value using a Consumer Price Index (CPI), which expresses income in 2014$.

Expected future annual earnings. To construct a measure of the persistent component of annual earnings, I rely on expected future annual earnings, which I measure as the answer to the question 'What do you believe your annual earnings will be in 4 months?' (question

[33] For instance, the same question presented in categorical terms for those who skipped L4 is 'Roughly speaking, what are your annual earnings, before taxes and other deductions, on your [current/main] job? Please include any bonuses, overtime pay, tips or commissions.' with proposed brackets as answers.

OO2e2). This question is also in the Labor Market module. As a validity check, I build an alternative measure of expected future earnings, based on other questions of the Labor Market module asking the respondents about the probability they assign to various income-related events between now and fourth months from now (e.g. receiving job offers, having their current employers match the offers).

I again assume that expected annual earnings in four months refers either to expected monthly earnings in four months extrapolated over the calendar year in which people will be in four months, or to the expected annual earnings over the calendar year in which people will be in four months. Under this latter assumption, since people are surveyed in November, their response captures expected annual earnings over the next calendar year. I show in Table 9 below that, consistent with this assumption (or with both these assumptions if earnings shocks mostly occur in January-March), the difference between annual earnings and expected annual earnings is 50% higher when the questions take place in November than when they take place in March and July (in logs and after detrending from demographics excluding quarters): people report much starker expected annual earnings change when the period they are asked about falls into a calendar year that is different from the current one. I deflate the value of expected future annual earnings using a quarterly CPI, which expresses expected future annual earnings in second quarter of 2014$. In general in the survey people are asked to 'not adjust for inflation and report annual earnings in today's current dollars' so I treat the reported expected annual earnings four months from a quarter t as given in quarter t dollars.

	March $t+1$ -Nov. t	Nov. t - July t	July t - March t
$ln(E_t[y^i_{t+1}]) - ln(y^i_t)$	0.013	0.007	0.009
Observations	2,117	2,596	2,170

Table 9: Covariance between log-persistent earnings and realized earnings growth

Probability to be employed next period. To construct a measure of persistent earnings, I also rely on the reported probability to be employed next period. I use the answer to the question 'What do you think is the percent chance that four months from now you will be [unemployed and looking for work or unemployed and not looking for work]' (question OO1) in the Labor Labor Market module.

Alternative measure of expected future annual earnings and variance of future annual earnings. I build an alternative measure of expected future annual earnings, in which I try to get rid of some possible response noise by using questions that force respondents to think about the probability of the main events that would modify persistently their earnings. This alternative measure exploits a set of questions, also in the Labor Survey module of the SCE, about the respondents' probability to receive job offers, the wage range of these offers, the probability to accept such offers, the probability that their current employer matches the offers they would receive, and the probability that they become non-employed, all four months from now. Specifically, the questions that I use are 'What do you think is the percent chance that within the coming four months, you will receive at least one job offer from another employer? Remember that a job offer is not necessarily a job you will accept.' (OO2e), 'Think about the job offers that you may receive within the coming four months. Roughly speaking, what do you think the annual salary for the best offer will be for the first year? Note the best offer is the offer you would be most likely to accept.' (OO2a2), 'If you were to receive a job offer from another employer at a higher salary, what do you believe is the percent chance your current employer will match the salary offer?' (OO2f), 'How much do you make before taxes and other deductions at your [main/current] job, on an annual basis? Please include any bonuses, overtime pay, tips or commissions ?' (L3, the question I use to measure current annual earnings), and 'What do you think is the percent chance that four months from now you will be [unemployed and looking for work, unemployed and NOT looking for work]?' (OO1). From these I build the different states of the world that each respondent faces, the earnings that he or she would receive in each state of the world, and the probability that he or she attach to the realization of each state of the world. More precisely, I build six possible states of the world in four months:

- people still have the same employer and they did not receive any job offer that their current employer would have been willing to match: their annual earnings is unchanged and equal to their current annual earnings;

- people still have the same employer but they did receive a job offer that their current employer has been willing to match: their annual earnings moves to 110% of their expected annual salary for the offer they would be most likely to accept;

- people have a different employer because they receive a job offer that they accepted: their annual earnings moves to their expected annual salary for the offer they would be most likely to accept;

- people have a different employer but not because they received a job offer: their annual earnings increases to 125% of their current annual earnings;

- people become self-employed: their annual earnings shifts down to 50% of their current annual earnings;

- people are unemployed or out of the label market: their annual earnings are zero.

My results are not very sensitive to the three assumptions I make, which are that in the 2nd, 4th and 5th scenarii, people's annual earnings are at 110%, 125% and 50% of their reported best offer or annual earnings (I tried for instance 100% 90% and 80% with similar results). The probabilities to have the same employer, to have a different employer, to be self-employed, to be unemployed or out of the labor market are reported, to receive at least one job offer, and that one's employer matches an external offer are directly reported by the respondents. I then recover the respondents' expected future income from this distribution. I also deflate the value using a quarterly Consumer Price Index and expressed the second quarter of 2014$.

I also use these expected values of earnings in each possible state of the world and the probabilities attached to them to compute a measure of the variance of future earnings that each individual faces.

Marginal propensity to consume. I build each measure of the MPC out of a transitory change in resources from the responses to two questions in the Household Spending module of the SCE. The first question is 'Now imagine that next year you were to find yourself with 10% less household income. What would you do?' (question QSP13new).[34] The second question, for those who would not use all of it on one thing thus whose MPC is not zero or one, asks them to quantify: 'Please indicate what share of the lost income you would cover by [Reducing spending, Reducing savings, Increasing borrowing].' (question QSP13a). I drop the 10 observations for whom the sum of these shares is not one. My baseline definition of the MPC is the share of the lost income that is covered up with reduced spending or increased debt. Similarly, I build the MPC out of a positive shock from the responses to the questions 'Suppose next year you were to find your household with 10% more income than you currently expect. What would you do with the extra income?' (question QSP12n) and 'Please indicate what share of the extra income you would use to [Save or invest, Spend or donate, Pay down debts]' (question QSP12a). I drop the 12 observations for whom the sum of these shares is not one. My baseline definition of the MPC corresponds to the share of the extra income that is used to spend and donate, or pay down debts. Since the question has 'next year' in it, I assume people answer on what they would do about with their spending/saving/debt over the next year—in addition, most of the previous questions about consumption and saving in this module ask people what they would do over the next twelve months, which further suggests people are framed to think over a twelve months horizon in this part of the survey.

The questions are about a one-time change in income next year, which is a transitory shock, but this transitory nature of the shock is not stressed in the question. For this reason, I check whether the answers are consistent with responses from other surveys asking about a hypothetical shock that is explicitly transitory. I find that, when comparable, the responses to my non-explicitly transitory shocks are similar to the responses to the explicitly transitory shocks presented in these other surveys. In Fuster, Kaplan, and Zafar 2020 the question is 'Now consider a hypothetical situation where you unexpectedly receive a one-time payment of [$500,$2,500,$5,000] today. We would like to know whether this extra income would cause

[34] The possible answers are [Cut spending by the whole amount, Not cut spending at all, but cut my savings by the whole amount, Not cut spending at all, but increase my debt by borrowing the whole amount, Cut spending by some and cut savings by some, Cut spending by some and increase debt by some, Cut savings by some and increase debt by some, Cut spending by some, cut savings by some and increase debt some].

you to change your spending behavior in any way over the next 3 months.', which includes the word 'one-time'. The options are then similar to those in the SCE: first people report whether what they would do, and then they are asked about the exact percentages. The average reported share of $2,500 and $5,000 income gain that would be used for spending (excluding debt repayment) over the next three months are 0.11 and 0.14 (Table A6). In my sample, the average reported share of a 10% annual income gain that would be used for spending (excluding debt repayment) over the next year is 0.15, thus very close—despite the remaining differences in the question. Because they do not report average fraction used for debt repayment I cannot compare with my data on this dimension. In Crossley, Fisher, Levell, and Low 2021, the wording is almost the same as in Fuster, Kaplan, and Zafar 2020. The share of a $500 gain that would be used for spending over the next three months is 0.11 (Table 1). This suggests that households do seem to interpret the question in the SCE in the same way they interpret a question about an explicitly transitory income shocks.

Importantly, I treat reported increased debt and repaid debt as changes in consumption. People reporting they would increase their debt in response to an income loss are rare, so what makes a difference is my treating repaid debt in response to an income gain as increased consumption. I argue there are compelling theoretical and empirical reasons to do so. Theoretically, since most debts have to be paid, the debts that are repaid with the extra income gain would probably have been repaid anyway over the following year. The question is: what would have people done to repay the debt if they had not received the extra income gain? Would they have cut their consumption or their savings? The literature on the response to an anticipated transitory income loss (which having to repay a pressing debt should be close to) suggests that people mostly cut their spending (see e.g. Gelman, Kariv, Shapiro, Silverman, and Tadelis 2018 on the response of expenditures to the 2013 U.S. government shutdown). Furthermore, liquidity constrained people who have no liquid savings and find it difficult (or too expensive) to borrow would mechanically have to cut their spending to repay their debt. These hand-to-mouth people alone constitute 25% to 40% of the population (see Kaplan, Violante, and Weidner 2014). Empirically, in my sample, considering increased debt and repaid debt as changes in consumption yields values that are much closer to the findings based on natural experiments. As reported in Table 11 Appendix B.4, the average yearly reported MPC of total consumption out of negative and positive shocks are 0.797 and 0.545 when treating debt variations as consumption variations, but 0.759 and 0.151 when not doing so. Now, the natural experiment literature finds that the yearly MPC of total consumption out of positive shocks is large: Parker, Souleles, Johnson, and McClelland 2013 find that more than 50% of the 2008 tax rebate is consumed over the next three months; taking into account the recent findings of Orchard, Ramey, and Wieland 2022 and Borusyak, Jaravel, and Spiess 2022 that MPCs measured with this methodology are overestimated and should be closer to 30% for total consumption (10% for strictly nondurables), so a reported MPC out of a positive shock of 54.5% over a year is consistent with an MPC of 30% over a quarter; Fagereng, Holm, and Natvik 2021 find that

70% of a lottery prize is consumed over the following year for lottery prizes between $5,200 and $8,300, which is the range closest to the shocks I examine here. A yearly MPC associated with total consumption of 15.1% is therefore below even the conservative estimates. In addition, excluding debt variations implies a very large asymmetry between the MPCs out of negative and positive shocks that might difficult to account for.

Wealth. In the main specification, I use a categorical measure of liquid (or semi-liquid) wealth. It is based on the answer to the question 'If you added up all the money in these accounts that you and your family members have invested in [Checking or savings accounts, Money market funds, CDs (Certificates of Deposit), Government/Municipal Bonds or Treasury Bills, Stocks or bonds in publicly held corporations, stock or bond mutual funds, or investment trusts], which category represents how much they would amount to? ' (question HQ17) in the Housing module. This excludes housing wealth. The respondent's wealth falls into 14 possible categories, from 'Less than $500' to '$1,000,000 or more'. The Household Finance module provides other, continuous measures of wealth, but I lose many observations when I use those. However, I rely on the Household Finance module to check the validity of my wealth measure, as discussed in Appendix B.3.

Demographics. I obtain demographic characteristics from the Core module of the survey. These include dummies for gender, age group, educational attainment, willingness to take risks, state of residence and number of family members. The age groups have a range of 5 years from 25-30 to 50-55. The educational attainment categories are people with only high-school education, those with some college education but who did not complete it, and those who completed college. My measure of the willingness to take risks comes from the answer to the question of the core SCE survey 'How would you rate your willingness to take risks in daily activities?' (QRA2). In controls, I also use period dummies, for the month-year of the interview.

Consumption. In one of my alternative specifications, I rely on consumption rather than on the hypothetical MPCs to capture the effect of persistent earnings on the consumption response to transitory shocks. I build consumption from a combination of questions from the the Spending module and the Housing module. Indeed, there is no direct question about the household's *level* of consumption expenditures in the SCE. However, the Spending module reports information about the share of their total monthly spending that the respondents' households allocate to different consumption categories in a typical month (housing, utilities, food, clothing, transportation, medical care, entertainment, education). The Housing module further reports information about the level of typical monthly spending on housing.[35] I thus recover the level of household's typical spending on different consumption categories with a proportionality rule, using

[35]It includes mortgage, rent, maintenance and home owner/renter insurance, which are the same categories that are considered when people are asked about the share of spending they devote to housing.

the level of housing spending, the share of total spending devoted to housing, and the share of total spending devoted to each of the other consumption categories. I do this for each category available (housing, utilities, food, clothing, transportation, medical care, entertainment and education), and sum them to obtain consumption. Note that the framing of the questions about the expenditures shares suggests that people would not include large purchases in the typical spending they report. Indeed, not only are they asked about typical spending, but just before this question about the allocation of their typical monthly spending among categories, the survey asks about large infrequent purchases (e.g. home appliances, computers, furniture, cars, trips, some of which hardly fall in any of the typical spending categories), so households likely exclude these large purchases from their answers to the next question about their typical spending. Importantly, this limitation to typical spending does not apply to the MPCs that are about total spending. My measure of yearly consumption is the typical monthly consumption spending multiplied by 12. Because this measure is based on multiple answers from different modules, I can only build it for a substantially restricted set of observations. I deflate this measure of consumption with a CPI index and express it in 2014$.

B.3 Validity of the wealth measure

To check the validity of my wealth measure, I compute its equivalent with variables from the Household Finance module, which contains more details but is observed for a smaller number of people in my sample. More precisely, I build from the Household Finance module a variable corresponding to the sum of the wealth from defined contribution accounts, the wealth from individual retirement accounts, and the wealth from checking and savings accounts, CDs, stocks, bonds, mutual funds, Treasury bonds. I also transform my categorical baseline wealth measure from the Housing module into a continuous one: I attribute to each respondent a wealth equal to the the lower bound of the wealth category they belong to—putting them at 0 when they answer ’less than $500.

	Baseline (made continuous)	Based on HF module
Mean	70,522	69,196
Standard dev.	171,207	225,292
P5	0	0
P25	5000	0
P75	50,000	30,000
P95	1,000,000	1,020,000
Observations	1,851	1,851

Table 10: Checking the consistency of wealth responses across questions

Table 10 presents the average and distribution characteristics of the two measures of wealth among people for whom both are observed. The first line shows that the averages are remarkably similar. The second line indicates that, as expected, since the baseline measure (based on a categorical variable that I transformed into a continuous one) only moves by thresholds, its variance in the sample is smaller. Finally, the third to sixth lines suggest that the distributions around the averages are not too dissimilar. The baseline measure has thinner tails, with both less households at zero wealth and less households at the highest levels of wealth. Overall, this confirms that, although my measure of wealth is obtained from a question in the Housing module, it captures the same information as the questions in the Household Finance module, and the answers are consistent across the two modules.

B.4 Summary Statistics

	Mean	Coef. of var.	Obs.
Annual earnings	63,471	(0.637)	1,117
Expected annual earnings (direct)	66,392	(0.636)	1,117
Expected annual earnings (indirect)	65,135	(0.637)	874
Expected proba to be empl.	0.976	(0.072)	1,117
MPC neg. (incl. increased debt)	0.797	(0.337)	1,097
MPC pos. (incl. repaid debt)	0.545	(0.648)	1,113
MPC neg. (excl. increased debt)	0.759	(0.387)	1,097
MPC pos. (excl. repaid debt)	0.151	(1.261)	1,113
Annual consumption (household-level)	60,951	(0.767)	1,007

Table 11: Descriptive statistics on the main variables

Descriptive statistics. Table 11 present descriptive statistics on the main variables for the respondents in the selected sample for whom I jointly observe current annual earnings, my measure of persistent earnings, at least one MPC out of a negative or a positive shock, and the consumption-related demographics (state of residence and family size). Current annual earnings is on average $63,471$ in 2014$. Expected future annual earnings is $66,392$, above people's reported current annual earnings: people expect on average some earnings growth over the next four months. The average of the alternative, indirect measure of expected annual earnings, based on questions about the probabilities of future income-related events is close, as it is equal to $64,466$. It confirms that the expected annual earnings that people report is consistent with their responses to questions about their probability of employment, job offers, and other income-related events. The reported probability of employment indicates that people mostly expect to still be employed four months from now.

The next four lines details the variables related to my measures of MPCs. As discussed when I present the MPC variables, the MPCs out of negative and positive transitory shocks are consistent with natural experiment results. In contrast, assuming that variations in debt have no subsequent effect on consumption, the MPC out of a positive shock gets very low (0.151), and substantially below what natural experiments suggest for the MPC of total consumption out of an income gain over the year following the income gain.

The last line presents the average value of annual household typical consumption. Comparing it with annual earnings (on the first line), it is a little smaller than the average annual earnings of the household head (or of one of the household heads). Considering that most respondents have more than one earner in the household, and might receive revenue other than earnings, a household's consumption expenditures should be a substantially smaller fraction of the household's income than of the earnings of the head.

	Sample share	Coef. of var.	Obs.
Female	0.524	(0.954)	1,117
Age below 25	0.045	(4.622)	1,117
Age between 25 and 30	0.149	(2.395)	1,117
Age between 30 and 35	0.150	(2.386)	1,117
Age between 35 and 40	0.170	(2.210)	1,117
Age between 40 and 45	0.152	(2.361)	1,117
Age between 45 and 50	0.164	(2.260)	1,117
Age between 50 and 55	0.171	(2.203)	1,117
Completed college	0.661	(0.717)	1,117
Some college	0.274	(1.629)	1,117
No college	0.065	(3.783)	1,117
Willingness to take risk of 1 (Not willing at all)	0.038	(5.061)	1,117
Willingness to take risk of 2	0.169	(2.217)	1,117
Willingness to take risk of 3	0.226	(1.849)	1,117
Willingness to take risk of 4	0.221	(1.878)	1,117
Willingness to take risk of 5 (Very willing)	0.206	(1.965)	1,117
One family member	0.169	(2.217)	1,117
Two family members	0.318	(1.466)	1,117
Three family members	0.201	(1.992)	1,117
Four family members	0.184	(2.110)	1,117
Five family members	0.084	(3.300)	1,117
Six family members	0.028	(5.921)	1,117
Seven family members or more	0.016	(7.817)	1,117
Surveyed in Nov 2015	0.235	(1.803)	1,117
Surveyed in Nov 2016	0.257	(1.701)	1,117
Surveyed in Nov 2017	0.269	(1.647)	1,117
Surveyed in Nov 2018	0.238	(1.789)	1,117

Table 12: Descriptive statistics on the main demographics

Table 12 presents some descriptive statistics on the demographic characteristics of the respondents, including only the respondents in the selected sample for whom I jointly observe my measure of persistent earnings, my categorical wealth variable, at least one MPC out of a negative or a positive shock, and the time-varying demographics (state of residence and family size). It shows that 52% of the respondents are female, roughly one third are below age 35, another third between age 35 and age 45, and a final third between age 45 and age 55. 66% of the respondents completed college, 27% have some some college education, and 7% no college education. Respondents declare themselves quite willing to take risks, with one fifth of them self-selecting into the highest category. 17% of the households are composed of one person, 32% of two persons, and the rest of third persons of more. Finally, the sample is quite equally distributed across the different survey years. I do not report the distribution across states of residence because of the large number of states, but it is not peculiar.

B.5 Robustness of the method to build persistent earnings

Using expectations over the next four moths. Given the earnings specification that I assume, and my objective of measuring the level persistent component, rather the shocks to this component over a given period, there is nothing a priori problematic in using expected annual earnings four months from now. The objective of the method is: (i) to still capture the persistent component of annual earnings; (ii) to get rid of the transitory component of annual earnings. If the persistent component is still present in what people expect four months from now, while the transitory component is no longer present, the method is valid. The results are not sensitive either to the probability of non-employment that I assume. More precisely, I use $p_{\nu_{t-1}} = 0.972$, that is the probability of still being employed four month from now. This is correct if most of the non-employment shocks occur between December and March and last for a year, so the probability of still being employed four months from now in November of a given year would approximately equal the probability to still be employed over the next year. However, I also try the more conservative assumption that the probability of being non-employed next year (conditional on being employed in November of the current year) is $p_{\nu_{t-1}} = 0.972^3 = 0.918$. Results are similar.

Timing of the transitory and persistent shocks. The method I use to identify persistent earnings does not depend on assumptions about the timing of the transitory and persistent shocks, that is, the moment when they occur within the year. It only requires that annual earnings be adequately represented by my specification (3.1)-(3.3)—so I have a variable that does not contain the current transitory component but contains the current persistent component (weighted by ρ)—, and many possible timings of the shocks are consistent with this general process. The variables η and ε may be functions of underlying shocks occurring at different points in the calendar year (not just at the beginning of the year) and several times per year (not just once), and still be consistent with (3.1)-(3.3).[36]

[36] For instance, imagine that at month m of year t, monthly earnings is $y_{t,m} = (1-\nu_t)e^{p_{t-1}}e^{\eta_{t,m}}e^{\varepsilon_{t,m}}e^{\alpha}e^{g(t)}$. The

I also show that there exist some specifications in which annual earnings are the sum of monthly earnings with shocks occurring at different points within the year that are consistent with the more restrictive subcase of (3.1)-(3.3) proposed by Guvenen, Karahan, Ozkan, and Song 2021, in which η and ε are drawn from mixtures of normal distributions (I exhibit an example of this in Appendix B.6).

Interpretation of the earnings-related questions. The question I use to build expected future annual earnings is 'What do you believe your annual earnings will be in 4 months?'. People might have different interpretation of what annual earnings in a given month is. My baseline interpretation is that they think about their expected monthly earnings at that given month and extrapolate over a calendar year (say monthly earnings have a monthly component, then extrapolating would be substituting this monthly component with that of all other months and summing to get an annual measure of earnings). Another possible interpretation is that people report their expected annual earnings over the calendar year of the given month in the future they are asked about. Both interpretations are consistent with my method. Furthermore, since the requirement of the method is simply to have a variable that does not incorporate the transitory component but does incorporate the persistent component (weighted by ρ), it is not very sensitive to the interpretation. In particular, the following interpretations of the question would barely affect the method:

- People declare as annual earnings in a given month in the future the sum of their expected monthly earnings in this month and the eleven months following: in that case a modest bias could arise because annual earnings would not be exactly proportional to $(e^{p_t^i})^{\rho}$. Only ten months out of twelve would be so (from March of year $t+1$ to December of year $t+1$) and the remaining two (January and February of year $t+2$ would be proportional to $(e^{p_t^i})^{\rho^2}$. Given that ρ is close to one, the bias should be small.
- People declare as annual earnings in a given month in the future their earnings conditionally on being employed, forgetting to take into account the scenario in which they would be unemployed. This means I should not divide expected future earnings by the probability to still be employed in three months to obtain expected earnings conditionally on being employed because people already report this conditional value. Alternatively, people may confuse income with earnings and treat unemployment benefits as earnings. This means that I should remove unemployment benefits times the probability to be non-employed from their expectations. I re-run the main analyses under these two assumptions and show in Appendix C.6 that the results are unaffected—this is not surprising since the probability to be unemployed in four months conditionally on being employed is very small (0.024).

shocks $\eta_{t,m}$ and $\varepsilon_{t,m}$ are monthly and can be drawn from an arbitrary distribution that can depend on the current month-year and on the individual's demographics. Then, denoting $e^{\varepsilon_t} = min_{m=1}^{12} e^{\varepsilon_{t,m}}$ and $e^{\eta_t} = \sum_{m=1}^{12} e^{\eta_{t,m}} e^{\varepsilon_{t,m}-\varepsilon_t}$, the sum $y_t = \sum_{m=1}^{12} y_{t,m}$ is consistent with the yearly specification (3.1)-(3.3)

Overestimation of the persistence ρ. The study of Rozsypal and Schlafmann 2017 suggests that, in their reported expectations, people are biased and overestimate the persistence of their earnings. Such an overestimation means that, when I build persistent earnings, the residual should be divided by $\rho* > \rho = 0.991$ if people erroneously expect the persistence to be $\rho*$, larger than the true persistence $\rho = 0.991$ estimated in administrative data by Guvenen, Karahan, Ozkan, and Song 2021. In that case, I estimate the effect of $perst^{\rho*/\rho}$ rather than the effect of $perst$. Since the authors still suggest $\rho* \leq 1$, the maximum bias would be to estimate $perst^{1/0.991}$, with $(1/0.991) = 1.009$.

Dependency between ρ and p_t^i. In the baseline specification, following Guvenen, Karahan, Ozkan, and Song 2021, I assume that the persistence ρ of the persistent component is a fixed parameter. However, the study of Arellano, Blundell, and Bonhomme 2017 suggests that the persistence is not fixed and depends on the value of the persistence component: $\rho = \rho_t^i = \rho(p_t^i)$. Since p_t^i is not a regressor in the linear regression I run to obtain ρp_t^i, this simply means that what I measure is not ρp_t^i but $\rho_t^i p_t^i$. When I divide the residual by 0.991, my measure of persistent earnings is no longer p_t^i but $(\rho_t^i/0.991)p_t^i$: it now also includes a term $\rho_t^i/0.991$ that correspond to the deviation of the persistence from 0.991. Thus what I measure is no longer just the persistent component of earnings but the persistent component of earnings adjusted for its persistence.

B.6 Microfounding the earnings process as the sum of monthly earnings

I examine under which conditions the yearly earnings specification of Guvenen, Karahan, Ozkan, and Song 2021 could write as a sum of monthly earnings, which is what people receive in practice. The specification I assume in Section 3 is more general than the specification of Guvenen, Karahan, Ozkan, and Song 2021 (because it does not make assumptions about the distributions from which the shocks are drown for instance), but I try to exhibit monthly processes that are consistent with the exact specification of Guvenen, Karahan, Ozkan, and Song 2021 because it has been shown to fit the data well, and it is the one I use in the numerical simulations. I identify two cases in which the yearly earnings assumed in Guvenen, Karahan, Ozkan, and Song 2021 can rewrite as the sum of monthly earnings.

The first case is straightforward: if the shocks occur once per year, at the beginning of the year, monthly earnings are the same every month $m \in [1:12]$ and equal to one twelfth of annual earnings. Their sum immediately writes up as the process proposed in Guvenen, Karahan, Ozkan, and Song 2021:

$$y_{t,m}^i = \frac{1}{12} y_t^i = \frac{1}{12}(1 - \nu_t^i)(e^{p_{t-1}^i})^\rho e^{\eta_t^i} e^{\varepsilon_t^i} e^{\alpha^i} e^{g(t)}$$

The second case allows persistent shocks to occur at different points in time over the year (or alternatively transitory shocks to occur at different points in time over the year). The other components are constant over the year. Thus, the monthly earnings only differ because of the monthly value of the persistent shock $e^{\eta^i_{t,m}}$:

$$y^i_{t,m} = \frac{1}{12}(1-\nu^i_t)(e^{p^i_{t-1}})^\rho \underbrace{e^{p^i_{t,m}}}_{\text{Monthly}} e^{\varepsilon^i_t} e^{\alpha^i} e^{g(t)}$$

With probability p_η, the persistent shock occurs at the beginning of the year and is drawn from a given distribution $\mathscr{S}^S$ (indexing it with S for soon). The monthly persistent shock is then:

$$e^{\eta^i_{t,m}} = s^i_t \;\; \forall m \in [1:12]$$

In contrast, with probability $(1-p_\eta)$, the persistent shock occurs later on, say on month m^η, and is drawn from a given distribution $\mathscr{L}^L$ (indexing it with L for late). The monthly persistent shock is then:

$$e^{\eta^i_{t,m}} = \begin{cases} 1 \;\; \forall m \in [1:m^\eta - 1] \\ l^i_t \;\; \forall m \in [m^\eta - 1 : 12] \end{cases}$$

Annual earnings are:

$$y^i_t = \sum_{m=1}^{12} y^i_{t,m} = \begin{cases} e^{\eta^i_{t-1}}(1-\nu^i_t)(e^{p^i_{t-1}})^\rho e^{\varepsilon^i_t} e^{\alpha^i} e^{g(t)} s^i_t \text{ with prob. } p_\eta \\ e^{\eta^i_{t-1}}(1-\nu^i_t)(e^{p^i_{t-1}})^\rho e^{\varepsilon^i_t} e^{\alpha^i} e^{g(t)}\left(\frac{m^\eta-1}{12} + \frac{13-m^\eta}{12} l^i_t\right) \text{ with prob. } (1-p_\eta), \end{cases}$$

with s^i_t drawn from a a lognormal distribution $\mathscr{S}^S = \text{Lognormal}(\mu_{\eta,1}, \sigma^2_{\eta,1})$, and $(\frac{m^\eta-1}{12} + \frac{13-m^\eta}{12} l^i_t)$ drawn from a log-normal distribution $\text{Lognormal}(\mu_{\eta,2}, \sigma^2_{\eta,2})$ (so l^i_t is drawn from a three-parameter log-normal $\mathscr{S}^L = \text{Lognormal}(-\frac{m^\eta - 1}{12}, \mu_{\eta,1} + ln(\frac{12}{13-m^\eta}), \sigma^2_{\eta,1}))$. This corresponds exactly to the annual earnings specification proposed by Guvenen, Karahan, Ozkan, and Song 2021.

B.7 Testing for anticipations

I examine the value of the covariance between the res^i_t, that is, the residual from a regression of the log of expected future annual earnings conditional on employment over demographics and time dummies, and the quasi-growth in log-earnings $\Delta^\rho ln(y_{t+1})$ among people who are employed both at t and at $t+1$. Indeed, according to my specification, that covariance is zero when future transitory shocks are ε_{t+1} are not anticipated, but strictly positive and equal to the variance of ε_{t+1} in the population when ε_{t+1} is anticipated at t:

$$cov_t(res_t, \Delta^{\rho} ln(y_{t+1})) = \begin{cases} cov_t(\rho p_t, \eta_{t+1} + \varepsilon_{t+1} - \rho\varepsilon_t + g(t+1) - \rho g(t) + (1-\rho)\alpha) \\ \quad = 0 \text{ without anticipation} \\ cov(\rho p_t + \varepsilon_{t+1}, \eta_{t+1} + \varepsilon_{t+1} - \rho\varepsilon_t + (1-\rho)\alpha + g(t+1) - \rho g(t)) \\ \quad = var_t(\varepsilon_{t+1}) > 0 \text{ with anticipation} \end{cases}$$

The covariance between the res_t^i and the quasi-growth in log-earnings $\Delta^{\rho} ln(y_{t+2})$ is also zero in the absence of anticipation but strictly positive and equal to $\rho^2 var_t(\varepsilon_{t+1})$ when ε_{t+1} is anticipated at t:

$$cov_t(res_t, \Delta^{\rho} ln(y_{t+2})) = \begin{cases} cov_t(\rho p_t, \eta_{t+2} + \varepsilon_{t+2} - \rho\varepsilon_{t+1} + (1-\rho)\alpha + g(t+2) - \rho g(t+1)) \\ \quad = 0 \text{ without anticipation} \\ cov_t(\rho p_t + \varepsilon_{t+1}, \eta_{t+2} + \varepsilon_{t+2} - \rho\varepsilon_{t+1} + (1-\rho)\alpha + g(t+2) - \rho g(t+1)) \\ \quad = \rho^2 var_t(\varepsilon_{t+1}) > 0 \text{ with anticipation} \end{cases}$$

I compute these covariance in the survey data, for the respondents in the selected sample for whom I jointly observe current annual earnings, my measure of persistent earnings, and at least one MPC out of a negative or a positive shock at t. In this exercise, because I work with earnings data from the Labor Market module, which takes place every four months, a period (the difference between t, $t+1$, and $t+2$) corresponds to four months.

	$cov_t(\frac{res_t^i}{\rho}, \Delta^{\rho} ln(y_{t+1}))$	$cov_t(\frac{res_t^i}{\rho}, \Delta^{\rho} ln(y_{t+2}))$
Value conditional on observing MPC and wealth	.004	.002
	(.006)	(.011)
Observations	877	278

Table 13: Covariance between log-persistent earnings and realized earnings growth

Table 13 presents the value of these covariances. They are small and not significantly different from zero. This suggests that my measure of persistent earnings, based on res_t, does not include any anticipated component at t of future shocks at $t+1$, because res_t does not covary with the realized earnings growth between t and $t+1$, nor with earnings growth between $t+1$ and $t+2$. Removing the minimum earnings threshold to include even those below the threshold and computing the covariance over this larger sample yield even smaller and less significant covariances.

B.8 Testing for independent additive shocks

The earnings specification I assume predicts that the individual variance of future earnings conditionally on being employed writes is approximately proportional to $(perst_t^i)^2$ (with the proportionality coefficient dependent on demographics and period dummies):

$$var_t^i(y_{t+1}^i) = \underbrace{\left((e^{p_t^i})^\rho e^{\overline{\varepsilon}} e^{\overline{\alpha}} e^{\overline{g}}\right)^2}_{\approx (perst_t^i)^2 \text{ for } \rho \approx 1} \underbrace{\left(e^{\alpha^i - \overline{\alpha}} e^{g(t+1)-\overline{g}}\right)^2 var_t^i(e^{\eta_{t+1}^i}) var_t^i(e^{\varepsilon_{t+1}^i - \varepsilon})}_{\text{Depends only on demographics and period dummies}}$$

Indeed, under my assumption is that the distributions of $e^{\varepsilon_{t+1}^i}$ and $e^{\eta_{t+1}^i}$ are the same conditional on demographics and period dummies, their variances do not depend persistent earnings but only on demographics and period dummies. On the contrary, when there exist independent additive shocks χ, the variance of future earnings is the sum of a term proportional to $(perst_t^i)^2$ and of the variance of χ_{t+1}^i.

I test this by regressing individual-level measures of the variance of future annual earnings conditional on being employed, detrended from the effect of demographics, over persistent earnings and persistent earnings squared allowing for a non-zero intercept. The way in which I detrend the variance is the same as the way in which I detrend the log of expected future annual earnings to build persistent earnings.[37] The specification is:

$$var_t^i(y_{t+1}^i|_{empl}) = a_1 + a_2 perst_t^i + a_3 (perst_t^i)^2 + \xi_t^i$$

	Variance over squared proba. of employment
Intercept	-6.213e+08
	(1.621e+09)
Persistent earnings	45937.98
	(55445.098)
Persistent earnings^2	0.791**
	(0.396)
R^2	0.719
Observations	862

Table 14: Effect of persistent earnings on the variance of future earnings

Table 14 presents the results. They are consistent with my earnings specification. The intercept a_1 is not statistically significant. It is relatively small as the value of $6.213e+08$ represents only 8% of a one standard deviation of $var_{t,i}(y_{t+1}^i|_{empl})$ (detrended) in the sample. The coefficient associated with the level of persistent earnings is not significant either. How-

[37] I regress its log over the same demographics used to build persistent earnings. I then take an exponential of the residual and normalize it so its average value is the same as the average value of the initial variance (in level).

ever, the coefficient associated with the square of persistent earnings is positive and significant. To get a sense of what the magnitude of this coefficient means, note that if the parameters were such that $\rho = 1$ (persistent shocks are permanent), $\alpha^i = \bar{\alpha}$ (no individual variability), and $var_t^i(e^{\eta_{t+1}^i}) = var_t^i(e^{\varepsilon_{t+1}^i - \varepsilon}) = 1$ (shocks are drawn from normalized distribution), then this coefficient should be equal to one. The fact that I estimate it to be 0.791 suggests that the parameters are different from this particular case, but not too far from it.

B.9 Comparison with Arellano et al (2021)

Although my sample is in the US while they study Spain, my results are not inconsistent with the finding of Felgueroso, García-Pérez, Jansen, and Troncoso-Ponce 2018 and Arellano, Bonhomme, Vera, Hospido, and Wei 2021 that 'inequality in income risk is related to the prevalence of high unemployment, but also to the large share of short-term temporary employment'.[38] I do two things: first, I show that I can reproduce the result of Arellano, Bonhomme, Vera, Hospido, and Wei 2021 that the coefficient of variation of future income is decreasing in current income; second, I show that I can also reproduce their result that that this is driven by people with low or zero earnings (not employed or with low attachment to the labor market) and that, among people who declare themselves as employed in the survey, the coefficient of variation of future earnings is not significantly related to current earnings, as would be the case with the specification that I rely on. Indeed, the income specification I assume implies the following coefficient of variation:

$$\begin{aligned}
CV_t^i &= \frac{\sqrt{var_t^i(y_{t+1}^i)}}{E_t^i[y_{t+1}^i]} \\
&= \frac{(e^{p_t^i})^\rho e^{\overline{\varepsilon}} e^{\alpha^i} e^{g(t+1)} sd_t^i(e^{\eta_{t+1}^i}) sd_t^i(e^{\varepsilon_{t+1}^i - \overline{\varepsilon}})(1 - p_{v_t^i})}{(e^{p_t^i})^\rho e^{\overline{\varepsilon}} e^{\alpha^i} e^{g(t+1)} E_t^i[e^{\eta_{t+1}^i}] E_t^i[e^{\varepsilon_{t+1}^i - \overline{\varepsilon}}](1 - p_{v_t^i})} \\
&= \frac{sd_t^i(e^{\eta_{t+1}^i}) sd_t^i(e^{\varepsilon_{t+1}^i - \overline{\varepsilon}})}{E_t^i[e^{\eta_{t+1}^i}] E_t^i[e^{\varepsilon_{t+1}^i - \overline{\varepsilon}}]}
\end{aligned}$$

By assumption, these standard deviations and expected values are independent of earnings because, conditional on demographics and period dummies and on the unemployment status, people draw shocks from the same distributions. Importantly, the probability of non-employment $p_{v_t^i}$ cancels out from the expression of the coefficient of variation. Therefore, the negative effect of persistent earnings on the probability of non-employment cannot explain the negative relation between the coefficient and persistent earnings that Arellano, Bonhomme, Vera, Hospido, and Wei 2021 document. However, what can explain it is that, my specification only implies the coefficient of variation is independent of earnings among employed people, while Arellano, Bonhomme, Vera, Hospido, and Wei 2021 include unemployed people and people

[38] See p.8 in Arellano, Bonhomme, Vera, Hospido, and Wei 2021.

with low attachment to the labor market, with low earnings and presumably a high coefficient of variation, in their exercise.

Method. I use two different methods to compute the effect of income (or earnings) on the coefficient of variation of future income (or earnings). With the first one, which I refer to as the 'group-level' method, I compute the coefficients of variations within a group of respondents that have the same set of predictors. This gives a sense of the risks faced by people with these characteristics. This is close to what Arellano, Bonhomme, Vera, Hospido, and Wei 2021 do. More precisely I measure the mean absolute deviation and mean of income (or earnings) among groups of individuals with the same demographic characteristics (the ones I use to build persistent earnings), same type of job (public, private for profit, non-profit, family business or other) and same job sector—either in which they are or were employed—at each period. For individuals within a group with the same characteristics, the coefficient of variation is the ratio of the mean absolute deviation of income (or earnings) within their group over the mean of income (or earnings) within their group. I regress people's coefficient of variation over their mean income (or earnings) at the previous quarter, that is, the mean income (or earnings) at the previous quarter within the group they belonged to at the previous quarter. The first column reports the coefficient of this regression. The sample is the people in the selected group for whom I observe current income (or earnings), current wealth, and at one least one current MPC.

With the second method, which I refer to as the 'individual-level' method, I use my individual-level measure of variance and of the expected value of future income (or earnings), based on questions about the probability that respondents attribute to different possible states of the world. I build the individual-level coefficient of variation of future income (or earnings) as the square-root of the individual-level variance of future income (or earnings) over the individual-level expected value of future income (or earnings). At the same time, I observe the current annual income (or earnings) of the respondents. I thus regress this current annual income (or earnings) over the individual-level coefficient of variation, controlling for the demographics I use to build persistent earnings and for period dummies. The second column reports the coefficient of this regression. The sample selection is the same as with the first method, but there are more missing observations.

Note that, in the Labor Market module, what I primarily observe is earnings. To measure income, I use the sum of earnings and unemployment benefits. I compute those benefits as the product of the respondents' reported earnings before unemployment with the unemployment insurance replacement ratio of the current quarter (obtained from the United States Department of Labor see `https://oui.doleta.gov/unemploy/repl_ratio/repl_ratio_rpt.asp`, such that the average in the sample is 0.460). When I examine income, I consider the coefficient of variation of future income, and I regress it over current income; when I examine earnings, I consider the coefficient of variation of future earnings and I regress it over current earnings.

CV	Group-level variance	Individual variance
Income, all	-2.83e-07**	-5.53e-07
	(1.37e-07)	(4.61e-07)
Observations	528	461
R^2	0.0081	0.0322
Earnings, all	-5.84e-07***	-3.81e-06***
	(2.04e-07)	(1.26e-06)
Observations	484	445
R^2	0.0167	0.0495
Earnings, employed	-3.99e-07	8.32e-08
	(2.53e-07)	(8.95e-08)
Observations	1,239	5,298
R^2	0.0167	0.0053

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

Table 15: Effect of persistent earnings on the variance of future earnings

Results. Table 15 presents these results. The first column of the first line shows that the finding of Arellano, Bonhomme, Vera, Hospido, and Wei 2021 is true in my survey data: the coefficients of variations built within groups are decreasing with the past income within groups. Incidentally, I also regress the coefficients of variation within groups over the current income within groups and the relation is negative as well. In the second column of the first line, the effect of income on the coefficient of variation of future income is still negative with the individual-level measures, but no longer significant.

The second line shows that this relation is even stronger for earnings when I still include both employed and non-employed people. This is not surprising since non-employed people have zero earnings, so their current earnings are presumably lower than their current income, while the coefficient of variation of their future earnings is larger.

The third line shows that the negative relationship disappears when the regressions are run only among employed people. This confirms the validity of my assumption that among employed people the coefficient of variation is independent of persistent earnings. However, the first two lines do suggest that unemployed people do not draw their shocks from the same distributions as others because their coefficient of variation is higher. This is the main reason why I focus on employed respondents, and select out the non-employed.

C Extra empirical results

C.1 Excluding people with strictly negative wealth

	MPC neg.	MPC pos.
Persistent earnings in \$10,000(a_2)	.017***	.016**
	(.005)	(.006)
Earnings in \$10,000 (a_3)	-.01*	-.011*
	(.005)	(.006)
One std. dev. persistent earnings	0.057	0.055
Average MPC	0.793	0.525
R^2	0.186	0.263
Observations	968	984

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

Table 16: Effect of persistent earnings on the MPC excluding people with strictly negative wealth

What the theory section predicts is that only people with positive wealth should respond more to transitory shocks when their persistent earnings are higher. The reason is that, for people with positive wealth, an increase in persistent earnings reduces the ratio of wealth to persistent earnings, while the opposite is true of people with negative wealth. As a result, if the theory is true, I should observe that, when I restrict the sample to people with positive wealth, the effect gets stronger. Table 16 presents the results from the estimation of the baseline specification (4.1) over a reduced sample that excludes respondents with strictly negative wealth. More precisely, I select only those for whom the continuous measure of total wealth obtained from the Household Finance module is either positive or unobserved—and select out those for whom it is strictly negative. The number of people with strictly negative wealth is eventually quite small. Table 16 shows that the point estimates get a little larger over this reduced sample—consistent with the theory that the effect is positive for those with positive wealth—but not significantly so.

C.2 Bootstrapped

	MPC neg.	MPC pos.
Persistent earnings in $10,000 ($a_2$)	.015***	.014*
	(.005)	(.00799903)
Earnings in $10,000 ($a_3$)	-.009*	-.009
	(.005)	(.006)
One std. dev. persistent earnings	0.048	0.046
Average MPC	0.797	0.545
R^2	0.166	0.236
Observations	1097	1113

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

Table 17: Effect of persistent earnings on the MPC

Table 17 presents the results when the standard errors are bootstrapped with 500 iterations. The bootstrapping loop includes both the first-stage building of the persistent earnings variables, and the second stage estimation of their effect on the MPCs. The effect remains significant at the 1% for the MPC out of a negative shock, and is significant at the 10% for the MPC out of a positive shock.

C.3 Total earnings only

	MPC neg.	MPC pos.
Earnings in $10,000	.001	-0.000
	(.003)	(.003)
Average MPC	0.798	0.546
R^2	0.157	0.23
Observations	1108	1125

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

Table 18: Effect of persistent earnings on the MPC

To understand what happens if one does not treat the persistent component of earnings separately from the rest of earnings, I estimate a specification in which I only consider total earnings. Table 18 presents the results. They show that the effect of total earnings is very small and no longer significant. The point estimate implies 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. This is in line with existing results on the impact of total earnings on people's response to transitory shocks, which find mostly non-significant effects (see e.g. Parker, Souleles, Johnson, and McClelland 2013, Boutros 2021, Parker, Schild, Erhard, and Johnson 2022).

C.4 Results controlling for transitory earnings instead of total earnings

	MPC neg.	MPC pos.
Persistent earnings in \$10,000 (a_2)	0.007***	0.006
	(0.003)	(0.004)
Earnings in \$10,000 (a_3)	0.006	-0.004
	(0.013)	(0.015)
One std. dev. persistent earnings	0.024	0.019
Average MPC	0.797	0.545
R^2	0.163	0.232
Observations	1097	1113

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

Table 19: Effect of persistent earnings on the MPC

I examine the effect of controlling for transitory earnings instead of controlling for total earnings. I build transitory earnings as the ratio of total earnings over persistent earnings. It therefore includes the deviations of the transitory component from its sample mean but also the deviations of the fixed effect component and the age trend component from their sample means, so this component is not purely transitory. This different control changes the interpretation of the coefficient associated with persistent earnings: in the baseline, the coefficient measures the effect of an increase in persistent earnings combined with a decrease in the transitory component of earnings (or rather in the non-persistent component) such that total earnings remains the same; here the coefficient measures the effect of an increase in persistent earnings, while the transitory component of earnings stays the same. However, this also changes the interpretation of the MPC: these are defined out of shocks that are proportional to income. Controlling for earnings means the MPCs are defined out of similar shocks, while controlling for transitory earnings means the MPCs of those with higher persistent earnings are defined out of larger shocks. Table 19 presents the results. The effect of persistent earnings on the MPCs is still positive but smaller. It is still significant for the MPC out of a negative shock but no longer significant for the MPC out of a positive shock. This is consistent with both the fact that the effect of persistent earnings is smaller when it is not combined with a decrease in transitory earnings, and the fact that, because the MPCs are defined out of shocks that are proportional to income, failing to control for total earnings reduces the effect of persistent earnings. These results can be interpreted as lower bounds of the effect of persistent earnings on the MPCs, while the baseline results can be interpreted as upper bounds.

C.5 More general specification

	MPC neg.	MPC pos.
Persistent earnings in $10,000	.013**	.021***
	(.006)	(.008)
Earnings in $10,000	-.013**	-.012*
	(.006)	(.007)
One std. dev. persistent earnings	0.044	0.07
Average MPC	0.797	0.545
R^2	0.193	0.266
Observations	1097	1113

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

Table 20: Effect of persistent earnings on the MPC with a general specification

I estimate a more general specification than the baseline one described by (4.1). Indeed, this baseline specification assumes that the MPC, that is, the partial effect of wealth on consumption, is a linear function of persistent earnings, total earnings, wealth dummies, and some controls. The more general specification allows for the MPC to be a quadratic function of persistent earnings and of total earnings, and for their effects to interact with each other and with the wealth dummies, as follows:

$$\begin{aligned} MPC_t^i = a_1 + a_2\, perst_t^i(1 + b_2 fam\, size_t^i + c_2 perst_t^i) + a_3\, earn_t^i(1 + b_3 fam\, size_t^i + c_3 perst_t^i + d_3 earn_t^i) \\ + a_4\, asset\, cat_t^i(1 + b_4 fam\, size_t^i + c_4 perst_t^i + d_4 earn_t^i) + a_5\, fam\, size_t^i + a_6 state_t^i + a_7\, dem_t^i + \xi_t^i, \end{aligned} \tag{C.1}$$

Table 20 presents the average effects of a change in persistent earnings and in total earnings implied by the estimates of this more general specification. The average effect of persistent earnings is still positive and significant. The point estimate of is a little smaller than in the baseline specification when looking at the MPC out of a negative shock while it is a little larger than in the baseline specification when looking at the MPC out of a positive shock. They are not significantly different from the baseline estimates. 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.

C.6 With alternative measures of persistent earnings

	MPC neg.	MPC pos.
Persistent earnings in \$10,000 (a_2)	0.018**	0.016*
	(0.007)	(0.008)
Earnings in \$10,000 (a_3)	-0.013*	-0.013*
	(0.007)	(0.007)
One std. dev. persistent earnings	0.060	0.053
Average MPC	0.79	0.533
R^2	0.196	0.269
Observations	866	880

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

Table 21: Effect of persistent earnings on the MPC

With persistent earnings based on an alternative measure of expectations. To understand whether people are consistent in their assessment of expected future annual earnings when asked directly about them without context versus when asked about the probability of events that would affect their future earnings, I develop an alternative measure of wealth (described in Appendix B.2). I then use this alternative measure of expected future earnings rather than my baseline, direct measure, to build the persistent component of earnings.

Table 21 presents estimates of the effect of this alternative measure of persistent earnings on the MPCs. The first line shows that the effect of persistent earnings is still positive and significant for both MPCs, and close to the point estimates I obtain in the baseline specification. The estimates are a little less precisely measured, probably due to the smaller sample size for which I observe this alternative measure of persistent earnings.[39] The second line shows that the effect of earnings, conditional on persistent earnings, is negative and mildly significant. The similarity in results confirms that the expected annual earnings that people report are consistent with what their other responses, not only with respect to their average level, but also in the way they vary with other variables.

[39]When I use my baseline measure of persistent earnings over the same reduced sample, the effect of persistent earnings on the MPC out of positive shock is even less precisely measured and only significant at the 15% level.

	MPC neg.	MPC pos.
Persistent earnings in \$10,000 (a_2)	.102***	.063**
	(.024)	(.03)
Earnings in \$10,000 (a_3)	-.018***	-.012*
	(.006)	(.007)
One std. dev. persistent earnings	0.088	0.054
Average MPC	0.798	0.548
R^2	0.177	0.233
Observations	1136	1152

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

Table 22: Effect of persistent earnings on the MPC when including also the influence of education and risk-taking on earnings in the persistent component

Including part of the α^i in the measure of persistent earnings. I modify the way in which I build persistent earnings, in order for the measure of persistent earnings to capture part of the individual fixed effect 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 reason I still remove the effect of the gender dummy is because, while the earnings are at the individual level, the MPCs are at the household level—so including the effect of gender on persistent earnings would mean including the effect of a variable that has a significant impact on persistent earnings but likely no impact on the MPC and would blur the relation between persistent earnings and the MPC. Because the effect of α^i does not fade away, I set the overall persistence of this measure of persistent earnings that includes both α^i and $e^{p_t^i}$ to $\rho = 1$. However, results are quasi identical when I set it to $\rho = 0.991$.

Table 22 presents the results. It shows that the effect of persistent earnings on the MPCs is still significant at the 1% and 5%. The point estimates are approximately six times larger than in the baseline case, at 0.102 and 0.063. However, the variance of persistent earnings is also larger with this definition, since it now includes the effect of individual-level demographics. As a result, the effect of a one-standard deviation change in persistent earnings is closer to the baseline: such a change raises the MPC out of a negative shock by 0.088 and the MPC out of a positive shock by 0.054 (instead of 0.048 and 0.046 in the baseline).

	MPC neg.	MPC pos.
Persistent earnings in \$10,000(a_2)	.015***	.017***
	(.005)	(.006)
Earnings in \$10,000 (a_3)	-.01*	-.011**
	(.005)	(.006)
One std. dev. persistent earnings	0.048	0.056
Average MPC	0.797	0.545
R^2	0.164	0.238
Observations	1098	1114

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

Table 23: Effect of persistent earnings on the MPC when people report expected earnings conditional on employment as expected earnings

	MPC neg.	MPC pos.
Persistent earnings in \$10,000(a_2)	.015***	.018***
	(.005)	(.006)
Earnings in \$10,000 (a_3)	-.01*	-.012**
	(.005)	(.006)
One std. dev. persistent earnings	0.049	0.057
Average MPC	0.797	0.545
R^2	0.165	0.238
Observations	1098	1114

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

Table 24: Effect of persistent earnings on the MPC when people report expected earnings plus UI as expected earnings

Misinterpretations of the expected earnings question. Tables 23 and 24 present the results under the assumption that people misinterpret the question about their annual earnings. In Table 23, they report their expected future annual earnings conditional on employment, forgetting to take into account the scenario in which they would be unemployed. In that case, I should not divide reported expected earnings by the probability to be employed. Table 23 shows that the results are almost unchanged when I do so. This is unsurprising since the probability to be employed in four months is very high and close to one in my sample of employed people.

In Table 24, people misinterpret what earnings are, and include the unemployment benefits they would get if they became unemployed in the expected value they report. In that case, I should divide reported expected earnings by the probability to be employed plus the probability to be unemployed times unemployment benefits (I describe exactly how I build unemployment benefits in Appendix B.9). Table 24 shows that, again, the results are almost unchanged when I do so. This is still unsurprising since the probability to become unemployed is low in my sample of employed people.

C.7 With MPC measures that exclude debt variations

	MPC neg.	MPC pos.
Persistent earnings in \$10,000 (a_2)	0.011**	-0.006*
	(0.006)	(0.003)
Earnings in \$10,000 (a_3)	-0.006	0.005*
	(0.005)	(0.003)
One std. dev. persistent earnings	0.038	-0.018
Average MPC	0.759	0.151
R^2	0.153	0.156
Observations	1097	1113

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

Table 25: Effect of persistent earnings on the MPC

I examine what happens when I rely on measures of the MPCs that assume changes in debt level have no impact on consumption later on and consider that only reported variations in spending constitute variations in consumption. This is in contrast to my baseline measure, in which I assume that when people state they use money to repay a debt, they then consume more than they would have had if they still had the debt to repay (so the debt repayment eventually corresponds to an increase in consumption), and when people state they increase borrowing, they then consume less than they would have had they not contracted a new debt (so the borrowing eventually corresponds to a decrease in consumption).

Table 25 presents the results of this estimation. It shows a stark difference between the responses to negative and positive shocks. In the first column, the effect of persistent earnings on the MPC out of a negative shock remains positive and significant at the 5% level. The point estimate is 0.011, a little smaller but similar to the baseline result. The effect of total earnings is negative, although the point estimate is no longer significant. The similarity with the baseline results is not surprising since only few people report that they would increase borrowing to cope with an income loss, so the baseline and alternative measures of the MPC out of a negative shock are not very different *(0.797 and 0.759 respectively).

In contrast, the effect of persistent earnings on the MPC out of a positive shock becomes negative when I exclude debt repayment from the MPC. This suggests that there exists a group of consumers with some debt to repay and a relatively high level of persistent earnings who report that they would use an income gain to repay their debt. As a result, when debt repayment is not considered consumption, the MPC of these consumers gets low and it reverses the sign of the relation between persistent earnings and the MPC. Importantly, the average MPC obtained with this alternative measure is very low, at 0.151, and much below natural experiments estimates of the yearly MPC of total consumption out of positive transitory income shocks (which is in general at least above 0.30 as discussed when I construct the MPC variable). It is therefore probable that, although people with some debt and a relatively high level of persistent earnings use an income gain to repay their debt, they then consume more than they would have had they not repaid their debt. This increased consumption would be visible in natural experiments but not in survey responses if people consider that they used the money to repay the debt—without reporting that having repaid the debt affects their consumption. This specification also presents some signs of misspecification: the effect of liquid wealth on the MPC is positive at some wealth levels, in contrast again to the findings from natural experiments and from the hypothetical MPC literature. This confirms the suspicion that the alternative MPC measure is not adequate.

C.8 Results using consumption rather than hypothetical MPCs

Statistical model. Because my measure of the MPC is based on a hypothetical question, it might be subject to some biases. I now consider a specification based on reported consumption rather than on these hypothetical questions. In that specification, 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. Indeed, the effect of non-housing wealth on consumption can measure a form of MPC, so the interaction would measure the effect of persistent earnings on this MPC. The specification that I estimate is:

$$\begin{aligned} Cons_t^i = \Big(& \tilde{a}_1 + \tilde{a}_2\, pers_t^i (1 + \tilde{b}_2 hh\, size_t^i) + \tilde{a}_3\, earn_t^i (1 + \tilde{b}_3 hh\, size_t^i) + \tilde{a}_4\, wealth_t^i (1 + \tilde{b}_4 hh\, size_t^i) \\ & + \tilde{a}_5\, hh\, size_t^i + \tilde{a}_6\, state_t^i + \tilde{a}_7\, fixed\ dem_t^i + \tilde{a}_8 period_t + \tilde{a}_9 age\ group_t^i \Big) \times wealth_t^i \\ & + \tilde{c}_1\, others_t^i + \tilde{\xi}_t^i. \end{aligned} \tag{C.2}$$

This specification is built to be consistent with the baseline specification of the MPC described by (4.1): differentiating both sides of this specification with respect to wealth yields exactly the baseline specification. The term $others_t^i$ is a vector of other determinants of consumption, unrelated to wealth.[40] The variable $wealth_t^i$ is based the same wealth category as I use in the main specification, transformed to be continuous: I set the wealth of the respondents equal to the lower bound of the wealth category they belong to—putting them at 0 when they answer 'less than $500. The reason why I convert the categorical variable into a continuous variable, is because otherwise I would have 14 interaction terms, and I am likely to loose some precision in my relatively small sample.

This specification is less robust than the previous one for several reasons. First, changes in wealth are not necessarily exogenous and might reflect a response to 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

[40]It includes the level of persistent earnings and of earnings, dummies for the year-quarter, number of family members, age category, state of residence, education level, willingness to take risk, and the level of housing wealth.

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. For these reasons, my preferred specification remains (4.1), which directly estimates the effect of persistent earnings on the MPC.

Implementation. I estimate (C.2) with a linear regression. The variable $perst_t^i$ is built as described in the previous subsection.

	Consumption
Persistent earnings in \$10,000 × Wealth ($\tilde{a}_2$)	0.027**
	(.013)
Earnings in \$10,000 × Wealth ($\tilde{a}_3$)	-0.039***
	(.01)
One std. dev. persistent earnings	0.089
Average effect of wealth	0.106
R^2	0.44
Observations	1014

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

Table 26: Effect of persistent earnings on the MPC

Effect of persistent earnings on the MPC. Table 26 presents selected results from the estimation of specification (C.2). The first line shows that a \$10,000 dollar increase in persistent earnings (holding total earnings, wealth and demographics constant) raises the partial effect of wealth on consumption by 0.027. This estimate is significant at the 5% level. The second line shows that a \$10,000 dollar increase in earnings (holding persistent earnings, wealth and demographics constant) reduces the partial effect of wealth on consumption by 0.039. This estimate is significant at the 10% level. Despite the limitations of this specification, the results are consistent with the baseline specification and with the theoretical prediction of the model: everything else being equal, at a higher level of persistent earnings, people are more sensitive to changes in wealth. The third line shows that a one standard deviation increase in persistent earnings raises the partial effect of wealth on consumption by 0.089. The average effect of wealth on consumption is smaller than the MPCs reported by people in the survey, probably due to the limitations I discussed. The point estimate is 0.106.

D Extra empirical results on the implications

D.1 Wealth and the MPC when controlling for persistent earnings only

	MPC neg.	MPC pos.
Less than \$500 assets		
\$500-\$999 assets	.047	.06
	(.047)	(.06)
\$1,000-\$1,999 assets	-.065	-.111*
	(.055)	(.062)
\$2,000-\$4,999 assets	.01	-.092*
	(.042)	(.054)
\$5,000-\$9,999 assets	.005	-.126**
	(.044)	(.055)
\$10,000-\$19,999 assets	-.011	-.244***
	(.044)	(.054)
\$20,000-\$29,999 assets	-.042	-.2***
	(.05)	(.059)
\$30,000-\$49,999 assets	-.063	-.263***
	(.047)	(.057)
\$50,000-\$99,999 assets	-.09*	-.32***
	(.046)	(.056)
\$100,000-\$249,999 assets	-.055	-.392***
	(.049)	(.056)
\$250,000-\$499,999 assets	-.157***	-.249***
	(.058)	(.075)
\$500,000-\$749,999 assets	-.075	-.513***
	(.085)	(.064)
\$750,000-\$999,999 assets	-.473***	-.672***
	(.077)	(.075)
More than \$1,000,000 assets	-.022	-.373***
	(.085)	(.099)
R^2	0.161	0.23
Observations	1097	1113

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

Table 27: Effect of wealth on the MPC with only persistent earnings controls

Table 27 reports the effect of changing wealth from 'Less than \$500' to a higher category when estimating a version of equation (4.1) that controls only for persistent earnings $pers_t^i$ and not for total earnings $earn_t^i$. The results are extremely similar to the ones obtained when controlling for both persistent earnings and total earnings (third and fourth columns of Table 3). This confirms that the source of the downward bias in the measure of the effect of wealth on the MPC is the failure to control for persistent earnings, not the failure to control for the other components of total earnings.

D.2 Gains from targeting fiscal stimulus among both employed and unemployed

	Average MPC neg.	Average MPC pos.
Income < 10th	.839	.575
Income < 25th	.811	.58
Income < 50th	.813	.583
Income < 75th	.804	.576
Income < 90th	.798	.566
All	.792	.554
Wealth < 13th	.852	.734
Wealth < 13th & Perst. > 50th	.858	.781

Table 28: Effect of income targeting on average MPCs among both employed and unemployed

	Average MPC neg.	Average MPC pos.
Earnings < 10th	.812	.541
Earnings < 25th	.813	.569
Earnings < 50th	.812	.581
Earnings < 75th	.804	.575
Earnings < 90th	.798	.566
All	.792	.554
Wealth < 13th	.852	.734
Wealth < 13th & Perst. > 50th	.858	.781

Table 29: Effect of earnings targeting on average MPCs among both employed and unemployed

Table 28 presents the average MPCs across income percentiles, rather than earnings percentile, and in a sample that include people who are not employed and people with annual earnings below 1,885$. Income is the sum of earnings and unemployment benefits, built as I detail in Appendix B.9. The results show that the difference between the average MPCs among people below the 10th income percentile and in the whole sample is more pronounced than when I consider earnings percentile and employed people only but it is still small.

Table 29 presents the average MPCs across earnings percentiles but in a sample that include people who are not employed and people with annual earnings below 1,885$. The difference between the average MPCs among people below the 10th income percentile and in the whole sample is a little more pronounced than among employed people only but still small.

E Simulations results

E.1 Calibration of the earnings process

	Value	Source		Value	Source
ρ	0.991	Guvenen et al. 2021	a_0	2.746	Guvenen et al. 2021
p_η	0.176	Guvenen et al. 2021	a_1	0.624	Guvenen et al. 2021
$\mu_{\eta,1}$	-0.524	Guvenen et al. 2021	a_2	0.167	Guvenen et al. 2021
$\sigma_{\eta,1}$	0.113	Guvenen et al. 2021	a_ν	-2.495	Guvenen et al. 2021
$\mu_{\eta,2}$	0.112	Guvenen et al. 2021	b_ν	-1.037	Guvenen et al. 2021
$\sigma_{\eta,2}$	0.046	Guvenen et al. 2021	c_ν	-5.051	Guvenen et al. 2021
σ_{p0}	0.450	Guvenen et al. 2021	d_ν	-1.087	Guvenen et al. 2021
p_ε	0.04	Guvenen et al. 2021			
$\mu_{\varepsilon,1}$	0.134	Guvenen et al. 2021			
$\sigma_{\varepsilon,1}$	0.762	Guvenen et al. 2021			
$\mu_{\varepsilon,2}$	-0.006	Guvenen et al. 2021			
$\sigma_{\varepsilon,2}$	0.055	Guvenen et al. 2021			
σ_α	0.472	Guvenen et al. 2021			

Table 30: Calibration of the earnings process

E.2 Examining the sources of the high average MPCs by removing elements from the baseline model.

	Simple earnings		Total wealth		Tot. w. & simp. earn.	
	neg.	pos.	neg.	pos.	neg.	pos.
Pers. earn. in \$10,000	0.159	0.687	0.010	0.036	-0.002	-0.002
Earn. in \$10,000	-0.177	-0.716	0.011	0.008	0.002	0.002
One std. dev. persistent earnings	0.379	1.632	0.039	0.139	-0.005	-0.005
Av. MPC	0.769	0.277	0.334	0.322	0.034	0.034
R^2	0.874	0.772	0.18	0.094	0.263	0.281
Observations	4893	4862	3440	3377	4813	4804
	No dem. trend		No UI nor transf.		Larger borr. lim.	
	neg.	pos.	neg.	pos.	neg.	pos.
Pers. earn. in \$10,000	0.014	0.011	0.011	0.008	0.020	0.015
Earn. in \$10,000	-0.001	-0.001	-0.003	-0.003	-0.013	-0.014
One std. dev. persistent earnings	0.056	0.042	0.042	0.030	0.078	0.060
Av. MPC	0.666	0.598	0.659	0.580	0.703	0.673
R^2	0.767	0.737	0.815	0.737	0.332	0.508
Observations	3477	3470	3478	3459	3458	3434

Table 31: Effect of persistent earnings on the MPC

The sources of the high average MPCs. In the simulated data, the MPCs are large, while large MPCs are notoriously difficult to generate in life-cycle models. I examine how removing certain features of the model, as detailed below, affects the MPCs, to determine which of these features are necessary to generate high MPCs. Table 31 shows that two necessary ingredients are the rich earnings process à la Guvenen, Karahan, Ozkan, and Song 2021, and the liquid wealth calibration.

Indeed, the first column of the first panel presents the results in simulations of a model that substitutes the rich earnings process for a simple transitory-persistent process. This specification generates counterfactual MPCs. While the average MPC out of a negative shock is 0.769, large and similar to its value in the baseline case, the average MPC out of a positive shock is only 0.277, which is smaller than in the survey data and smaller than what natural experiment suggests (for the yearly MPC of total consumption).

The second column of the first panel presents the results in simulations of a model that substitute the calibration of the discount factor to match the average liquid wealth in the survey data, for a calibration of the discount factor to match the average total wealth in the survey data. This specification generates MPCs that are both too small: the average MPCs are only 0.334 and 0.332.

The third column of the first panel shows that, removing both the rich earnings process and the liquid wealth calibration reduces starkly the average MPCs, which become both very small (0.034 and 0.034). This is consistent with the now well-know fact that life-cycle models, which usually incorporate only a simple income specification, yield small MPCs when the discount factor is calibrated to match total wealth: people with a relatively small earnings risk and a lot of wealth are not sensitive to transitory earnings shocks.

The second panel shows that removing either the demographic trend after age 49, the presence of a consumption threshold and of a minimum income threshold, or extending the borrowing limit has little impact on the average MPCs: they remain relatively large in these three cases. These elements might still interact with others, but their removal alone does not reduce the large MPCs that the model produces.

Simple transitory-persistent earnings specification. This version of the model assumes that earnings evolve according to the simple transitory-persistent earnings specification that is commonly used in simulations. With this specification, there are no individual fixed effects (the value of α^i is the same for everybody), no state of unemployment, no quadratic trend, and shocks are drawn from normal distributions rather than mixtures of normal distributions. The distributions are centered around zero and their variances are the variances of the most probable distribution in the baseline model: the variance of the transitory shock ε is 0.055, and the variance of the persistent shock η is 0.046. The common component α is chosen so the mean earnings are the same as in the baseline model ($ 58,251). The value of β that matches an average liquid wealth close to $3,561 in this version of the model is 0.977. The sample is bigger for this specification because everybody is employed while in the baseline model a fraction of people are selected out because not employed.

Total wealth calibration. This version of the model calibrates the discount factor so that it matches the amount of total wealth in the data, rather than just the amount of liquid wealth.

I compute the amount of total wealth in the data as the sum of all assets (including housing), minus all debt (including mortgages), from questions in the Household Finance module of the SCE. The outcome is deflated to be expressed in 2014$, as I do with all other variables.[41] My resulting total wealth average is $207,762. The value of β that matches an average wealth close to this amount is 1.031, which means people need to be very patient to hold such a large level of wealth.

Total wealth calibration & simple earning specification This version of the model combines the use of the simple earnings specification with a calibration of β that matches total wealth rather than liquid wealth. The value of β that I obtain is 1.002. The sample is also bigger for this specification because, with the simple earnings process, everybody is employed.

No demographics trend This version of the model drops the assumption that, after age 49, the discount factor gets multiplied by 0.985 because of demographic changes ($e^{\Delta\delta_t z_{i,t}} = 0.985$ if t≥49 and $e^{\Delta\delta_t z_{i,t}} = 1$ if t<49). The new assumption is that there are no demographic changes ($e^{\Delta\delta_t z_{i,t}} = 1$ for all t). The reason why I assume the change in demographics in the baseline is because of empirical papers documenting a change in consumption driven by demographics over the life-cycle, but this assumption is not standard in numerical simulations: this is why I verify that it is not driving the results. The value of β that matches an average liquid wealth close to $3,561 in this version of the model is 0.939.

No threshold nor UI This version of the model drops the assumption that the isoelastic utility function only applies to the consumption expenditures above a threshold of $2,175, and the assumption that people receive transfers bringing their income to $2,175 whenever their earnings

[41] The questions I use are 'Approximately what is the total current value of your [and your spouse's/partner's] savings and investments (such as checking and savings accounts, CDs, stocks, bonds, mutual funds, Treasury bonds), excluding those in retirement accounts?' (D16new) for liquid wealth, 'Considering all the Defined Contribution plans you and your spouse/partner may currently have with a current or previous employer, approximately, what is the total amount of money currently in all these account(s)?' (C2new) for wealth in defined contribution plans, 'Considering next any self-owned retirement accounts (IRAs), you and your spouse/partner may currently have, approximately, what is the total amount of money currently in all these account(s)?' (C2newx1) for IRA wealth, 'Now, think about the value of any additional savings or assets (such as cash value in a life insurance policy, a valuable collection for investment purposes, or rights in a trust or estate) you and your spouse/partner may have. Taken together, what do you think is the total value of these additional savings or assets (minus anything you owe on them)?' (D20e) for other non-housing wealth, 'About how much do you think your home would sell for on today's market?' (F4) and 'And what about your other homes(s)? About how much in total do you think the other home(s) you own would sell for on today's market?' (F14) for housing wealth, 'Approximately, what is the total amount of outstanding loans against your home(s), including all mortgages and home equity loans?' (F18) for housingrelated debt, and 'Next consider all outstanding debt you [and your spouse/partner] have, including balances on credit cards (including retail cards), auto loans, student loans, other personal loans, as well as medical or legal bills, but excluding all housing related debt (such as mortgages, home equity lines of credit, home equity loans). Approximately, what is the total amount of your [and your spouse's/partner's] current outstanding debt?', for other debt. 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. This coefficient is the same I use for adjusting liquid wealth when I do a liquid wealth calibration.

are below this level. Another way to say it is that I bring the thresholds for consumption and for receiving transfers to zero.

Large borrowing limit. This version of the model lets people have up to \$ 30,000 in debt, instead of only \$ 3,261. The value of β that matches an average liquid wealth close to \$3,561 in this version of the model is 0.981.