How Does Permanent Income Affect Consumption? Theoretical and Empirical Results on the Concavity of Consumption in Permanent Income

Jeanne Commault*

February 2025

Abstract

I prove analytically that, in the standard life-cycle model used throughout the macroeconomic literature, consumption is not linear but concave in the permanent income component of a standard transitory-income process. The reason for the concavity in income is the opposite of that behind the seminal result of concavity in (accumulated) wealth, established in Carroll and Kimball (1996). An increase in permanent income strengthens the need for precautionary saving, and more so at a high level of permanent income: consumption responds less to an increase in permanent income at high levels because the need for saving increases more. In contrast, an increase in wealth reduces the need for precautionary saving, but less so at a high level of wealth: consumption also responds less to an increase in wealth at a high level of wealth but now because the need for saving decreases less. I also find empirical evidence suggesting that the results holds in US survey data.

Key words: Marginal Propensity to Consume, Heterogeneity, Permanent Income

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

1 Introduction

What characteristics affect people's consumption response to income shocks? Both the Great Recession, which exposed the limits of the representative agent model to explain even aggregate phenomena, and the rise in income inequality documented over the past decades, which calls for a better understanding of how income heterogeneity translates into heterogeneity in people's behavior, have marked a renewed interest in the sources and the extent of heterogeneity in consumption behavior.

However, the main source of heterogeneity studied in the literature is wealth. Theoretically, the seminal paper of Carroll and Kimball 1996 shows that consumption is concave in wealth, which means that at lower levels of wealth, consumption is more responsive to a change in wealth or equivalently to a transitory income shock. This result served as a basis for later analyses of the Great Recession (e.g. Mian and Sufi). Empirically, from the early estimations of marginal propensities to consume to more recent measures based on detailed datasets.

The picture is very different when one considers heterogeneity in consumption behavior across levels of permanent income, that is, across the fixed effect component of people's income. Theoretical results focus on the cases in which consumption is linear in permanent income, that is, in which a change in permanent income has the same impact on consumption at all levels of permanent income: the permanent income hypothesis of Friedman shows that, in the absence of uncertainty or with quadratic preferences, consumption is linear in human capital, which scales in permanent income for a standard transitory-permanent income process; Carroll 2016 shows that the household's maximization problem scales in permanent income; Straub 2019 examines a particular case in which consumption is linear in a certain definition of permanent income: when permanent income is defined as a fixed-individual specific productivity and when initial assets scale in this productivity across people. Empirically, results specifically aimed at measuring concavity are rare. There exist indirect evidence, through the observation that estimates of the elasticity of consumption to a permanent income shock are below one. However, these results only imply concavity under the assumption that, in levels, consumption is proportional to permanent income or to permanent income raised to certain power that might not be one.

In this paper, I make two contributions to the study of heterogeneity in consumption behavior across levels of permanent income: (i) I prove that, in a standard life-cycle model with uncertainty and a transitory-permanent income process, consumption is strictly concave in permanent income; this is because permanent income increasingly strengthens people's precautionary motive; (ii) in a reduced form regression, I measure the permanent component of income and examine its effect on consumption; my result suggests more concavity than previously found.

The framework I consider is a standard life-cycle model in which households seek to maximize their expected lifetime utility, subject to a budget constraint and a terminal condition on wealth. Their utility is isoelastic. They face uncertainty because the income they receive at each period is stochastic. In the framework that I first consider, income is a transitory-permanent process, the product of a permanent component that evolves as a multiplicative random walk and of a transitory component that is a non-serially correlated shock. In the generalization, I let income be a flexible function of different components, independent of each other.

First, I show that precautionary saving is increasing and convex in the permanent component of the transitory-permanent process. Intuitively, everything else being equal, because shocks to future income are multiplied by current permanent income, an increase in permanent income increases the need for precautionary saving, and—because of the shape of the marginal utility—all the more so when the increase in permanent income is applied to an already high level of permanent income. This implies that consumption is concave in permanent income: at a higher level of permanent income, an increase in permanent income raises the optimal level of saving more, and consumers save a larger share of the permanent income gain than they would at a lower level of permanent income. I also show that consumption is concave in the transitory component of the transitory-permanent process. The mechanism is the converse: everything else being equal, an increase in transitory income raises the share of total resources that is certain thus reduces the need for precautionary saving. Because of the shape of marginal utility, this is all the more true at an initially low level of transitory income. This implies that consumption is concave in transitory income: at a higher level of transitory income, and increase in transitory income does not reduce the need for saving as much, and consumers save a larger share of the transitory income gain than they would at a lower level of transitory income.[1]

Second, I examine the relation between permanent income and consumption empirically. To build a direct empirical measure of permanent income, I rely on the expected value of future income. Indeed, I show that, assuming the same transitory-persistent specification that Guvenen, Karahan, Ozkan, and Song 2021 find to match well US administrative data on earnings, the persistent component coincides with expected future income conditional on being still employed, detrended from the effect of future

[1]Note that although the concavity in transitory income is less new than the concavity in permanent income, since Carroll and Kimball 1996 establish the concavity of consumption in assets (with a different method than mine), I keep it to show the behavior of precautionary saving and the parallel with the permanent income case.

demographic variables, and adjusted for persistent income not fully persisting from the current period to the next. Intuitively, a change in current persistent income affects both current and future expected income while a change in transitory income affects current income but not expected future income. That is why expected future income can be used to disentangle the two.

I implement this method in the Survey of Consumer Expectations (SCE) run by the Federal Reserve Bank of New-York, which reports both expected future income and current consumption. As a robustness check, I build two measures of expected future income: one from a direct question about it, and one from a series of questions about the probabilities of future earnings-related events (job loss, job offers at different wage ranges, matching of the offer by one's employer). I also verify that my measure of current permanent income does not erroneously captures future shocks that people already know about and anticipate.

I find that consumption is significantly increasing and concave in permanent income. The estimates imply that, on average in the sample, a one dollar gain in permanent income raises nondurable consumption by 0.28 dollars on average, but this response is 0.37 at the 10th percentile of permanent income and only 0.17 at the 90th percentile. Such a variation corresponds to a more pronounced concavity than found in the existing literature. With a decomposition, I find the difference comes both from my relying on more general a level-specification, in which there are two parameters that govern the effect of permanent income, rather than on a log-specification, and from my relying on a direct measure of permanent income. I additionally show that both failing to account for nonlinearities in the effect of persistent income on consumption and imposing a specification in which a single parameter governs the level of the effect and concavity in persistent income lead to underestimating the average MPCP. The former also mechanically misses the heterogeneity in MPCPs, the latter underestimates it when one imposes that log-consumption is linear in the log of the total earnings of the household (while the true specification would be consumption being quadratic in persistent income).

2 Concavity

2.1 Model

Consumers' maximization problem At each period t, a consumer chooses its current consumption and the distribution of its future consumption as the solution of the following intertemporal optimization problem:

$$\max_{c_t,\ldots c_T} \quad \sum_{s=0}^{T-t} \beta^{t+s} E_t\left[u(c_{t+s})\right] e^{\delta_{t+s} z_{t+s}} \tag{2.1}$$

$$s.t. \quad a_{t+k+1} = (1+r_{t+k})a_{t+k} - c_{t+k} + y_{t+k} \quad \forall\, 0 \leq k \leq T-t, \tag{2.2}$$

$$c_{t+k} > 0, \quad \forall\, 0 \leq k \leq T-t, \tag{2.3}$$

$$a_T \geq 0. \tag{2.4}$$

The consumer is finite-lived with T the length of its life. It has time-separable preferences, and at each period t it derives utility from its contemporaneous consumption expenditures c_t. The period utility function $u(c)$ is isoelastic so its functional form is $u(c) = \frac{c^{1-\rho}-1}{1-\rho}$. Future utility is discounted by the factor β and depends on the effect of current deterministic characteristics z through the term $e^{\delta_{t+s} z_{t+s}}$. At each period, the household earns the stochastic amount y_t. The budget constraints (2.2) state that to store its wealth from one period to another the consumer only has access to a risk-free 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 certain interest rate r_t. The conditions (2.3) restrict consumption to be strictly positive at each period. They do not actually bind when utility is isoelastic. The terminal condition on wealth (2.4) states that the household cannot die with a strictly positive level of debt.

Income process The level of income is a flexible function of independent components: the component e^{p_t} which I am interested in, another component μ_t (possibly transitory but also possibly another permanent component independent of p_t or a combination of both), and demographic characteristics z_t:

$$y_t = f(e^{p_t}, e^{\mu_t}, z_t), \tag{2.5}$$

$$\text{with } e^{p_t} = g(e^{p_{t-1}}, \eta_t). \tag{2.6}$$

The term η denotes an independent shock that is serially uncorrelated. Its variance at each period is strictly positive, ruling out perfect certainty about future income. I assume that the component e^p is such that:

(i) income is increasing and concave in it:
$\frac{dy_t}{de^{p_t}} > 0$ and $\frac{d^2 y_t}{d^2 (e^{p_t})^2} \leq 0$

(ii) the component is increasing and concave in its past value:
$\frac{de^{p_t}}{de^{p_{t-1}}} \geq 0$ and $\frac{d^2 e^{p_t}}{d^2 (e^{p_{t-1}})^2} \leq 0$

These two assumptions hold for the permanent component of a standard transitory-permanent process.[2].[3] Note that this specification allows the value of e^{p-1} to influence the distribution at $t-1$ of η_t, in the spirit of Arellano, Blundell, and Bonhomme 2017.

First order condition The first order condition of the problem is:

$$u'(c_t) = E_t[u'(c_{t+1})]\underbrace{\beta(1+r)e^{\Delta\delta_{t+1}z_{t+1}}}_{R_{t+1}}$$

It states that, to maximize their total expected utility, consumers seek to equalize their current and future marginal utility of consumption, weighted by a deterministic term R_{t+1}.

2.2 Concavity

Theorem In the model described above consumption, for any component p_t that verifies assumptions (i) and (ii), for all $t < T$:

(a) $\frac{(1+r)^{T-t+1}}{\sum_{s=0}^{T-t}(1+r)^s} < \frac{dc_t}{da_t} < (1+r)$ and $\frac{1}{1+r}\frac{\partial y_t}{\partial e^{p_t}}\frac{\partial c_t}{\partial a_t} \leq \frac{dc_t}{dp_t} \leq \left(\sum_{s=0}^{T}\frac{1}{(1+r)^{s+1}}E_t\left[\frac{\partial e^{p_{t+s}}}{\partial e^{p_t}}\frac{\partial y_{t+s}}{\partial e^{p_{t+s}}}\right]\right)\frac{\partial c_t}{\partial a_t}$

(b) $\left(\frac{d^2c_t}{da_t e^{p_t}}\right)^2 < \frac{d^2c_t}{da_t^2}\frac{d^2c_t}{d(e^{p_t})^2}$

(c) $\frac{d^2c_t}{da_t^2} < 0$ and $\frac{dc_t}{d(e^{p_t})^2} < 0$

Note that although the first part of proposition (c) is the object of the paper of Carroll and Kimball 1996, I restate it because I rely on a different proving method that then helps understand the concavity in permanent income. The proofs are detailed in the Online Appendix and I propose here an intuition for each Proposition.

The intuition for the first part of Proposition (a) is that the marginal propensity to consume (MPC) out of a change in assets cannot be larger than (1+r) because this would mean that people would consume than what they initially gained, so their future consumption would fall while their current consumption would rise, and the change in current and future expected marginal utility would not be equal violating the Euler equation. The MPC out of a change in assets has to be larger than $\frac{(1+r)^{T-t+1}}{\sum_{s=0}^{T-t}(1+r)^s}$ because, when utility is isoelastic, a change in future consumption (in all states of the world)

[2] Incidentally, they also hold for the transitory component in a special case with $\frac{de^{p_t}}{de^{p_{t-1}}} = 0$ when the transitory component is an i.i.d. shock.

[3] Consider the process $y_t = e^{p_t}e^{\mu_t}$ with $e^{p_t} = (e^{p_{t-1}})^\rho e^{\eta_t}$, $0 < \rho \leq 1$, and $e^{\mu_t} = e^{\varepsilon_t + \theta_1\varepsilon_{t-1} + \ldots + \theta_q\varepsilon_{t-q}}$. It is true that $\frac{dy_t}{de^{p_t}} = e^{\mu_t} > 0$, $\frac{d^2y_t}{d(e^{p_t})^2} = 0$, $\frac{de^{p_t}}{de^{p_{t-1}}} = \rho(e^{p_{t-1}})^{(\rho-1)}e^{\eta_t} > 0$, $\frac{d^2e^{p_t}}{d(e^{p_{t-1}})^2} = \rho(\rho-1)(e^{p_{t-1}})^{(\rho-2)}e^{\eta_t} \leq 0$, and $var_t(y_{t+s}) = (e^{p_t})^2 var_t(e^{\eta_{t+1}}\ldots e^{\eta_{t+s}}) \underset{e^{p_t}\to 0}{\longrightarrow} 0$.

reduces expected future marginal utility more than a change in current consumption reduces current marginal utility. As a result, one can express the response of current consumption as larger than a deterministic, increasing function of the response of future consumption. Iterating backward from the last period T at which the response of consumption is $(1+r)$, I obtain this expression.

The second part of Proposition (a) states that the MPC out of e^p must be larger than the MPC out of assets multiplied by the effect of e^p on current income. This is because a rise in e^p raises the resources available for consumption at least as much as a gain in assets of magnitude $\frac{1}{1+r}\frac{\partial y_t}{\partial e^{p_t}}$, since it also raises future expected income. Thus, it must raise consumption at least as much as a gain in assets of magnitude $\frac{1}{1+r}\frac{\partial y_t}{\partial e^{p_t}}$ (and vice-versa for a loss). Also, the response to a change in permanent income is smaller than $\left(\sum_{s=0}^{T}\frac{1}{(1+r)^s}E_t\left[\frac{\partial e^{p_{t+s}}}{\partial e^{p_t}}\frac{\partial y_{t+s}}{\partial e^{p_{t+s}}}\right]\right)\frac{\partial c_t}{\partial a_t}$. Intuitively, because a gain e^p still increases uncertainty,[4] the response of consumption has to be smaller than what it would be if the change in resources that it generates was certain and equal to its expected value, the sum of current and future expected impact on income: $\left(\sum_{s=0}^{T}\frac{1}{(1+r)^s}E_t\left[\frac{\partial e^{p_{t+s}}}{\partial e^{p_t}}\frac{\partial y_{t+s}}{\partial e^{p_{t+s}}}\right]\right)$. Thus it must raise consumption less than a gain of assets of this magnitude.

Now, the proof of Proposition (b) is established using Cauchy-Schwarz inequality. I show that $\left(\frac{d^2c_t}{da_t e^{p_t}}\right)^2$ rewrites as a sum of elements whose squared values are smaller than the products of elements in $\frac{d^2c_t}{da_t^2}$ and $\frac{d^2c_t}{d(e^{p_t})^2}$.

The proof of Proposition (c) is based on the same Euler equation reasoning that I employ to establish Proposition (a). I begin with the case of concavity in assets. Deriving twice each side of the Euler equation, the second derivative of current consumption with respect to current assets is the sum of two things: a term that takes the same sign as the second derivative of future consumption with respect to future assets, and a term that captures the fact that the same change in current and in future consumption does not change the current and future expected marginal utility of consumption by the same amount because of uncertainty. I show that this second term is negative when utility is isoelastic: in that case, the fact that future consumption reduces the expected marginal utility more, which calls for current consumption to respond more and for precautionary saving to decreases, is less pronounced around higher levels of assets, so the reduction in precautionary saving caused by a gain in assets is smaller around higher levels of assets. By backward induction, because at the last period T consumption is linear in assets, and because if consumption at $t+1$ is concave in assets then consumption at t

[4]Technically, a rise in e^p comes generates a negative covariance between future income and the slope of the marginal utility of future consumption that reduces the amount by which the expected marginal utility of future consumption falls with future consumption, and thus the need for current consumption to respond more.

must be strictly so, consumption is strictly concave in all $t < T$.

The reasoning is similar for concavity in a component of income e^{p_t} that is independent of other components and verifies assumptions (i)-(ii), except that while a change in current assets only affects future consumption by influencing future assets, a change in the current value of this component affects future consumption by influencing both future assets and future permanent income. As a result, deriving twice both sides of the Euler equation, the second derivative of current consumption with respect to current permanent income is more complex: it is the sum of a term that depends on $\frac{d^2c_t}{d^2a_t}$, $\frac{d^2c_t}{da_te^{p_t}}$, and $\frac{d^2c_t}{d(e^{p_t})^2}$, plus a second term that captures the fact that the same change in current and in future consumption does not change the current and future expected marginal utility of consumption by the same amount because of uncertainty. Using Proposition (b) and assumption (ii), I show that the first term is negative when future consumption is concave in future assets and in future permanent income. The second term is also negative when utility is isoelastic.

2.3 Interpretation of concavity in terms of precautionary behavior

One direct way to understand what happens to people's behavior is to look at how precautionary saving evolves. Precautionary saving is defined as the difference between the level of consumption that people choose (in the presence of uncertainty), and the level of consumption that they would choose under perfect foresight, that is, in the absence of uncertainty, if their future income was certain and equal to its expected value, everything else being equal. We know from Kimball 1990 that consumption is strictly smaller in the presence of uncertainty than in its perfect foresight counterpart (denoted with an index PF). Thus the difference, precautionary saving denoted PS, is strictly positive:

$$PS_t = c_t^{PF} - c_t > 0 \tag{2.7}$$

Permanent income and precautionary saving I consider a component of income that, besides assumptions (i)-(ii), verifies a third one:

(iii) the uncertainty about future income approaches zero as permanent income approaches zero: $var_t(y_{t+s}) \underset{e^{p_t}\to 0}{\longrightarrow} 0$

Note that the permanent component of a standard transitory-permanent income process verifies this assumption. Under this assumption, when e^{p_t} approaches zero, the variance of future income approaches zero, thus the consumer maximization problem approaches its perfect foresight counterpart. As a result, the response to a change in e^{p_t} also approaches its perfect foresight counterpart:

$$\lim_{e^{p_t} \to 0} \frac{dc_t}{de^{p_t}} = \frac{dc_t^{PF}}{de^{p_t}} \tag{2.8}$$

Because consumption is strictly concave in permanent income, $\frac{dc_t}{de^{p_t}}$ is strictly decreasing in permanent income. It means that for any $e^{p_t} > 0$ $\frac{dc_t}{de^{p_t}} < \lim_{e^{p_t} \to 0} \frac{dc_t}{de^{p_t}} = fracdc_t^{PF} de^{p_t}$. Note that because under perfect foresight consumption is linear in permanent income, $\frac{dc_t^{PF}}{de^{p_t}}$ is constant and is the same at all levels of assets as it is when e^{p_t} approaches 0. As a result:

$$\frac{dc_t}{de^{p_t}} < \frac{dc_t^{PF}}{de^{p_t}} \tag{2.9}$$

$$0 < \frac{d(c_t^{PF} - c_t)}{de^{p_t}} = \frac{dPS_t}{de^{p_t}} \tag{2.10}$$

Thus, when assumption (iii) holds, the fact that consumption increases concavely with e^{p_t} comes with precautionary saving increasing convexly with e^{p_t} (the convexity is straightforward as $\frac{d^2 PS_t}{de^{p_t 2}} = \frac{d^2 c_t^{PF}}{de^{p_t 2}} - \frac{d^2 c_t}{de^{p_t 2}} = -\frac{d^2 c_t}{de^{p_t 2}} > 0$). Precautionary saving increases with e^{p_t}, that is, consumption responds less to a change in e^{p_t} than it would under perfect foresight, because a gain in e^{p_t} increases the uncertain part of their total resources thus increasing their precautionary needs. This effects intensifies at higher levels of e^{p_t}: each additional gain in e^{p_t} raises precautionary needs more than the previous gain. Figure 1 depicts the mechanism graphically.

Transitory income (or assets) and precautionary saving Again, although the case of assets has been examined, I present it here as a point of comparison with the case of a component of income that verifies (iii). Because in the case of a purely transitory component of income (or of a change in assets) denoted e^{p^t}, $\frac{de^{p^t_{t+1}}}{de^{p^t_t}} = 0$, and the variance of future income is unaffected by a change in the value of this component. When the value of this component approaches infinity, the consumer maximization problem approaches its perfect foresight counterpart, because current, certain, resources approach the entirety of the resources that the consumers have to finance their consumption:

$$\lim_{e^{p^t_t} \to +\infty} \frac{dc_t}{de^{p^t_t}} = \frac{dc_t^{PF}}{de^{p^t_t}} \tag{2.11}$$

Because consumption is strictly concave in $e^{p^t_t}$, $\frac{dc_t}{de^{p^t_t}}$ is strictly decreasing in assets. As a result, for any $e^{p^t_t} < \infty$ $\frac{dc_t}{de^{p^t_t}}$ must be strictly larger than its limit value when a_t approaches

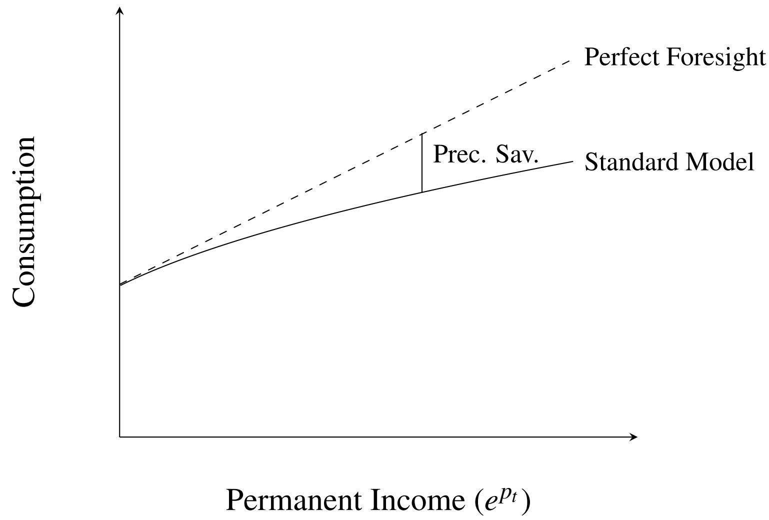

Figure 1: Graphical depiction of the behavior of precautionary saving with e^{p_t}

∞: $\frac{dc_t}{de^{p_t^t}} > \lim_{e^{p_t^t} \to +\infty} \frac{dc_t}{de^{p_t^t}} = \frac{dc_t^{PF}}{de^{p_t^t}}$. Because under perfect foresight consumption is linear in assets, $\frac{dc_t^{PF}}{da_t}$ is constant and is the same at all levels of assets as it is when a_t approaches ∞. As a result:

$$\frac{dc_t}{d\frac{dc_t^{PF}}{de^{p_t^t}}} > \frac{dc_t^{PF}}{d\frac{dc_t^{PF}}{de^{p_t^t}}} \tag{2.12}$$

$$\frac{d(c_t^{PF} - c_t)}{d\frac{dc_t^{PF}}{de^{p_t^t}}} = \frac{dPS_t}{d\frac{dc_t^{PF}}{de^{p_t^t}}} < 0 \tag{2.13}$$

Thus, when $e^{p_t^t}$ has no impact on the distribution of future income, the fact that consumption increases concavely with $e^{p_t^t}$ comes with precautionary saving decreasing convexly with $e^{p_t^t}$ (the convexity is straightforward as $\frac{d^2PS_t}{de^{p_t^t}{}^2} = \frac{d^2c_t^{PF}}{de^{p_t^t}{}^2} - \frac{d^2c_t}{de^{p_t^t}{}^2} = -\frac{d^2c_t}{de^{p_t^t}{}^2} > 0$). Precautionary saving decreases with e^{p_t}, that is, consumption responds more to a change in $e^{p_t^t}$ than it would under perfect foresight because a gain in this component of income increases their current, certain resources without affecting their future resources. This effects weakens at higher levels of $e^{p_t^t}$: each additional gain in $e^{p_t^t}$ reduces precautionary needs less than the previous gain. Figure 2 depicts the mechanism graphically.

2.4 Comparison with the linearity results

Th existing literature has established that the household's maximization problem was scaling in permanent income (Carroll 2016) and that in the particular case with permanent income being fixed over the life-cycle and initial assets being proportional to

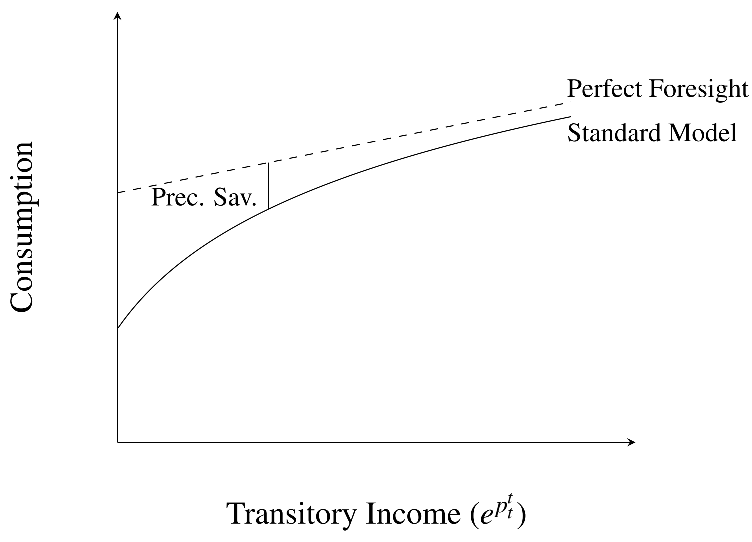

Figure 2: Graphical depiction of the behavior of precautionary saving with $e^{p_t^t}$

permanent income in the population, then consumption is linear in permanent income in the population (Straub 2019). More precisely, in this latter case, the linearity means a joint increase in permanent income and assets (so that assets remains the same proportion of permanent income) has the same impact at all levels of permanent income, conditional on the proportion of assets to permanent income.

I note that these results are not incompatible with mine. Indeed, Straub 2019's definition of permanent income is compatible with initial permanent income in a standard transitory-permanent income model. Although the effect of permanent income on consumption is independent of the level of permanent income when conditioning on the share of assets to permanent income, but this effect decreases with permanent income when conditioning on the level of assets.

3 Empirical investigation on income

3.1 Detrended expected future income as permanent income

I build an empirical counterpart to the permanent component of income in a transitory-permanent process. This observed variable is future expected income, detrended from the effect of demographic characteristics. Indeed, when income is a transitory-permanent process:

$$y_{t+1} = e^{p_{t+1}} e_{\varepsilon_{t+1}} e^{\delta_{t+1} z_{t+1}}, \text{ with } p_{t+1} = p_t + \eta_{t+1} \text{ and } \varepsilon_{t+1} \text{ independently drawn} \tag{3.1}$$

$$E_t[y_{t+1}] = e^{p_t} \underbrace{E_t[e^{\eta_{t+1}}] E_t[e^{\varepsilon_{t+1}}]}_{=1} e^{\delta_{t+1} z_{t+1}}. \tag{3.2}$$

The shocks are typically normalized such that $E_t[e^{\eta_{t+1}}] = E_t[e^{\varepsilon_{t+1}}] = 1$ (sometimes such that their log is zero, but I do a different normalization here). I then run a regression of $ln(E_t[y_{t+1}])$ on a set of demographic characteristics z. The residual from this regression, raised to the exponential and multiplied by the average value of $e^{\delta_{t+1} z_{t+1}}$ in the sample, is my measure of permanent income.

I then examine empirically whether $E_t[y_{t+1}]$ verifies assumptions (i), (ii) and (ii), conditional on demographic characteristics. Note that the existing evidence favoring the use of the standard transitory-permanent income process is that the empirical values of the autocovariance of income is consistent with such a process, and inconsistent with competing models. The three assumptions that I make are not specifically examined.

3.2 The Survey of Consumer Expectations (SCE)

To do so, I rely on panel data from the Survey of Consumer Expectations (SCE) of the Federal Reserve Bank of New-York, in particular because the surveys reports detailed information about people's expectations about their income. The panel is rolling and individuals remain in the survey for a maximum of a year. The questions about income expectations are in the Labor Market module of the survey, which participants answer every four months. This questions about future expectations that I use are only present in the SCE only since March 2015, and the data available goes up to March 2019, so my baseline sample covers March 2015-March 2019.

For employed individuals, I compute expected future income from their reported probabilities to receive job offers of different annual wage ranges, their probabilities to accept such offers at each wage range, the probability that their current employer matches an offers they would receive, and the probability that they become unemployed, all in four months from now. The value is deflated using a quarterly Consumer Price Index. The values are in the second quarter of 2014 \$. To build a detrended counterpart of expected future income, I regress it on a set of demographic characteristics: dummies for the date (quarter-year), the region, the number of children in the household, the age of the individual, its gender, and its experience at his or her current position, the industry sector in which he or she works, and the type of job (government, non-profit, private). Note that I consider two different options. In the first one, I detrend $E_t[y_{t+1}]$ from the effect of demographics at t. This makes it possible to keep more observations. In the second one, I detrend $E_t[y_{t+1}]$ from the effect of demographics at $t+1$, which is more consistent with the theory of expression (3.2), but means discarding more observations. Apart from the experience dummy, these characteristics are recorded in the Core part of the SCE, and not in the Labor Module Survey.

I build current income from a direct question about people's average annual income. The value is also deflated and converted in 2014 $.

I build people's expected variance of their future income from the same questions about the probability to receive job offers and to become unemployed, considering these different possibilities as different states of the world.

Finally, I build consumption from a combination of questions in the Spending module and in the Housing module. Indeed, a problem is that there is no direct question about the household's level of consumption expenditures. In the Spending module, there are, however, questions about the share of the household's total spending allocated to different types of consumption in a typical month (housing, utilities, food, clothing, transportation, medical care, entertainment, education and others). The framing of the question suggests excluding large purchases from this typical spending. In the Housing module, there are questions about the typical monthly spending on the components of the housing category included in the Spending module question.[5] I thus compute the level of household's total spending using a proportionality rule, from the level of housing spending and the share of housing in total spending. My measure of nondurable consumption spending includes utilities, food, clothing, transportation, and entertainment. My measure of nondurable and others consumption spending additionally includes the other spending category. Average yearly consumption is the typical monthly consumption multiplied by 12. Because the Housing module is only run once a year, it leaves me with a substantially restricted set of observations when I use consumption variables.

The set of demographics that I use for control in the regression is the same as the one I use for building the residuals.

I trim the top and bottom 5% of all these variables, replacing their values with non-reported (so the order in which I trim the variables does not matter).

Table 1 presents descriptive statistics on the mean and standard deviation of the variables I use. For consistency, I include only observations for which my measure of permanent income is observed. The first part of the table show that expected future

[5] This description is 'including mortgage, rent, maintenance and home owner/renter insurance'.

	Mean	Standard-Deviation
Current Inc.	56,316	31,838
Expected fut. inc. (before det.)	56,034	31,585
Expected fut. inc. (detrended)	55,988	31,761
Obs.		2,196
Nondurable cons.	25,511	17,148
Nondurable + other	29,812	20,402
Total cons.	53,952	33,214
Obs.		547

Table 1: Descriptive Statistics

income is a little lower and a little less volatile than current income. This is consistent with the idea that expected future income does not incorporate the transitory component of future income, making it less volatile. The values of expected future income before and after detrending are very similar, suggesting that demographics do not play a large role in determining their value. Also, the small difference between their mean comes from the fact that the mean value of demographics by which I multiply the residual from my regression is computed on a slightly different sample (in this table I consider observations for which income is non-zero). The fact that detrended expected future income is more volative than its non-detrended counterpart indicates that the shocks bring more volatility than changes in demographics. The second part presents household level consumption. Thus, note that the fact that annual current income is close to the annual consumption does not mean that an average households consume all their income: current income is computed as the individual level, while consumption is at the household level, and there is often more than one earning individual in a household. Nondurable consumption represents approximately half of total consumption (excluding large one-time purchases as discussed). The category of 'other consumption' is not small and quite volatile, which is why I exclude it from my baseline measure.

3.3 Testing assumptions (i)-(iii)

Assumption (i): income is increasing and linear or concave in permanent income
The specification that I estimate is:

$$y_t = \alpha_0 + \alpha_1 perm_t + \alpha_2 perm_t \times perm_t + \beta_1 dem_t \times perm_t + \xi_{t+1} \tag{3.3}$$

Testing assumption (i) corresponds to testing whether the marginal effect of $perm_t$ (through α_1, α_2 and β_1) is strictly positive, and whether its second order effect α_2 is non-significant or negative.

Future dem. in residual	no	yes
Constant	485	1252
	(609)	(810)
Perm. Inc.	0.997***	0.978***
(average effect)	(0.006)	(0.008)
Perm. Inc.^{2}	-3.48e-07*	-2.62e-07
(average effect)	(2.10e-07)	(2.50e-07)
R^2	0.960	0.961
Obs.	2,862	1,435
p-value RESET test	0.503	0.612

Table 2: Testing assumption (i) among employed individuals

Table 2 shows that income is increasing and linear (or mildly concave) in my measure of permanent income. Assumption (i) is not rejected. More precisely, Column 1 presents the results from estimating 3.3 detrending future expected income from the effect of current demographics rather than future demographics, to keep more observations. On average, an increase of one dollar in past permanent income leads to an increase of 0.997 dollar in current permanent income. This effect is mildly decreasing across levels of permanent income: a further one-dollar increase in permanent income reduces this response by $-3.48e-07$. The constant is not significant, suggesting that current income is proportional to permanent income. Column 2 presents the results from estimating 3.3 detrending future expected income from future demographics. The results are very similar to those in Column 1 although the second-order term is no longer significant. I run a specification test to check that the assumption is tested with a correct specification. The p-value of the RESET test is large, suggesting that non-linear combinations of the fitted value have no predictive power beyond that of the fitted value.

Assumption (ii): permanent income is increasing and linear or concave in past permanent income The specification that I estimate is:

$$perm_t = \alpha_0 + \alpha_1 perm_{t-1} + \alpha_2 perm_{t-1} \times perm_{t-1} + \beta_1 dem_t \times perm_{t-1} + \xi_t \quad (3.4)$$

Testing assumption (ii) corresponds to testing whether the marginal effect of $perm_{t-1}$ (through α_1, α_2 and β_1) is strictly positive, and whether its second order effect α_2 is non-significant or negative.

Table 3 shows that income is increasing and linear in my measure of permanent income. Assumption (ii) is not rejected. More precisely, Column 1 presents the results

Future dem. in residual	no	yes
Constant	4691**	3247
	(2008)	(1972)
Perm. Inc.	0.929***	0.953***
(average effect)	(0.020)	(0.023)
Perm. Inc.^{2}	6.17e-08	1.84e-07
(average effect)	(3.81e-07)	(4.72e-07)
R^2	0.887	0.919
Obs.	981	337
p-value RESET test	0.017	0.651

Table 3: Testing assumption (ii) among employed individuals

from estimating (3.4) detrending future expected income from the effect of current demographics rather than future demographics to build permanent income, to keep more observations. On average, an increase of one dollar in past permanent income leads to an increase of 0.929 dollar in current permanent income. This effect is constant across levels of permanent income: a further one-dollar increase in permanent income does not significantly change this response. The constant is mildly significant, suggesting that permanent income is not entirely proportional to past permanent income. Column 2 presents the results from estimating (3.4) detrending future expected income from future demographics to build permanent income. The results are very similar to those in Column 1 although the constant term is smaller and no longer significant. I run the same specification check. The RESET test indicates some misspecification with the estimation conducted in Column 1, suggesting that the reliable results are those of Column 2.

Assumption (iii): the variance of future income is proportional to permanent income The specification that I estimate is:

$$var_t(y_{t+1}) = \alpha_0 + \alpha_1 perm_t + \alpha_2 perm_t \times perm_t \quad (3.5)$$

$$+ \beta_2 dem_t \times perm_t \times perm_t + \xi_{t+1} \quad (3.6)$$

Testing assumption (iii) corresponds to testing whether the constant is zero, $\alpha_0 = 0$, which makes $var_t(y_{t+1})$ zero at $perm_t = 0$.

As the constant is not significant, these results are consistent with people's expected variance of their future income being proportional to their permanent income, thus approaching zero when permanent income approaches zero. More precisely, in Column 1, the expected variance of future income is a function of its permanent income (although

Future dem. in residual	no	yes
Constant	-1.73e+07	3.84e+07
	$(5.33e+07)$	$(8.32e+07)$
Perm. Inc.	5046**	3570
(average effect)	(2200)	(3280)
Perm. Inc.^{2}	1.059***	1.046***
(average effect)	(0.018)	(0.025)
R^2	0.887	0.980
Obs.	2,923	1,472
p-value RESET test	0.455	0.386

Table 4: Testing assumption (iii) among employed individuals

the effect is small in comparison to the value of the variance) and of the square of its permanent income. In Column 2, the variance of future income is entirely proportional to the square of permanent income, which is consistent with future income being proportional to permanent income. Th RESET test rejects a misspecification that would take the form of omitting combinations of the fitted value as predictors.

4 Concavity of consumption in permanent income

The specification that I consider is:

$$c_t = \alpha_0 + \alpha_1 perm_t + \alpha_2 perm_t \times perm_t + \beta_0 dem_t + \beta_1 dem_t \times perm_t + \xi_{t+1} \tag{4.1}$$

Table 5 shows evidence of concavity of consumption in my measure of permanent income. Column 1 presents the results of the estimation (4.1) when consumption is nondurable consumption. It shows that, at zero permanent income, an extra $ of permanent income raises nondurable consumption by 0.339, but this effect significantly decreases by $-1.09e-06$ with each additional $ of permanent income. Overall, the average effect of permanent income on nondurable consumption is 0.215. It decreases from 0.293 at the 10th percentile of permanent income to 0.119 at the 90th percentile of permanent income. The RESET test does not suggest that higher order values of current predictors are omitted. Column 2 presents the results of the estimation (4.1) when consumption is nondurable consumption plus the category of other expenditures. I do this because this category is sometimes included as nondurables in the literature. The results are similar. The concavity is a little more pronounced as the effect of permanent income on consumption decreases more steeply with permanent income, shifting from 0.384 at

Goods	Nondur.	Nondur. + oth.	Nondur.
Incl. assets	no	no	yes
Perm.Inc.	0.339***	0.426***	0.224***
	(0.063)	(0.070)	(0.083)
Perm.Inc.^{2}	-1.09e-06**	-1.28e-06***	-5.48e-07
	(4.45e-07)	(4.94e-07)	(6.01e-07)
Average effect	0.215***	0.280***	0.162***
	(0.021)	(0.026)	(0.027)
Effect at p(10)	0.293***	0.372***	0.201***
	(0.052)	(0.051)	(0.060)
Effect at p(90)	0.119***	0.167***	0.116**
	(0.035)	(0.041)	(0.046)
R^2	0.120	0.149	0.104
Obs.	668	666	365
p-value RESET test	0.941	0.894	0.938

Table 5: Measuring concavity of consumption among employed individuals

the 10th percentile to 0.116 at the 90th percentile of permanent income. Finally, Column 3 includes controls for liquid assets and its squared value the estimation. Because assets is obtained from the Finance module of the SCE and observed only for a subset of respondent, the number of observations drops to 365. Over this reduced sample, the effect of permanent income squared is no longer significant, but its sign is still negative.

4.1 Comparison with the literature

Existing results on the concavity of consumption in permanent income come from two types of measures. First, some studies rely on some observed persistent shocks and compute the response of consumption. Among them, Cochrane 1991 finds that an episode of illness of more than 100 days leads to a 11% to 14% decrease in consumption. An involuntary job loss leads to a 24% to 27% decrease in consumption. As size of the income loss implied by the shock is unclear, the exact elasticity and MPC out of permanent income remain unclear, but it is consistent with an elasticity below one, that is consistent with concavity when consumption scales in permanent income (possibly raised to a power different from one).

Second, semi-structural papers have been estimating the elasticity of consumption to permanent income shocks for a long time. Among them, the seminal paper of Blundell, Pistaferri, and Preston 2008 estimates this elasticity to be 0.64. Later studies, including those of Chatterjee, Morley and Singh (2017) and Crawley 2018, refine the analysis and suggest a slightly lower elasticity. Using a similar model of consumption but a different definition of permanent income and a different identification strategy, Straub 2019 finds the elasticity to be between 0.40 and 0.74. For an individual with such an elasticity whose nondurable consumption is half its permanent income (as is the case on average from 1), the MPC out of permanent income would be half the elasticity, thus comprised between 0.20 and 0.37.

This comparisons suggests that my results are within the range of the existing literature, although on the lower end of the set of existing results. I recomputed the results in log, as the existing literature does, and I find an elasticity of 0.41 (for nondurables) and 0.44 (including other durables), consistent with my level estimates.

However, although the level effects are on the lower end of existing estimates, estimating the concavity with two coefficients yields a stronger concavity than when concavity is a by-product of having an elasticity below one.

5 Conclusion

In this paper, I show that consumption is concave in any independent component of income such that income is increasing and linear or concave in it, and that is increasing and linear or concave in its past value. This encompasses the permanent component of income in the standard transitory-permanent process. In this particular case, the standard-deviation of future income scales in current permanent income. As a result, the concavity of consumption implies that precautionary saving is increasing and convex in permanent income: each additional gain of permanent income raises the need for precautionary saving, and more so at higher levels of permanent income. The opposite is true with a component of income that does not affect the variance of future income, such as transitory income. In that case, precautionary saving is decreasing and convex in transitory income: each additional gain of transitory income reduces the need for precautionary saving, but less so at higher levels of transitory income.

Using survey data from the SCE, I build an empirical counterpart to the permanent component of income in the standard transitory-permanent process. I show that my empirical counterparts verifies the assumption that ensure concavity (when consumption is the solution of a life-cycle model). I measure the concavity of consumption using this direct measure of permanent income. I find the concavity to be more pronounced than in the previous literature.

References

Arellano, M., Blundell, R., and Bonhomme, S. (2017). "Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework." *Econometrica* 85, pp. 693–734.

Blundell, R., Pistaferri, L., and Preston, I. (2008). "Consumption Inequality and Partial Insurance." *American Economic Review* 98.5, pp. 1887–1921.
Carroll, C. (2016). "Theoretical Foundations of Buffer Stock Saving." mimeo.
Carroll, C. and Kimball, M. (1996). "On the Concavity of the Consumption Function." *Econometrica* 64.4, pp. 981–92.
Cochrane, J. (1991). "A Simple Test of Consumption Insurance." *Journal of Political Economy* 99.5, pp. 957–76.
Crawley, E. (2018). "Time Aggregation in Panel Data on Income and Consumption." working paper.
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.
Kimball, M. (1990). "Precautionary Saving in the Small and in the Large." *Econometrica* 58.1, pp. 53–73.
Straub, L. (2019). "Consumption, Savings, and the Distribution of Permanent Income." *mimeo*.