Consumption Does Not Evolve as a Random Walk Around Small Income Shocks

Jeanne Commault*

February 2018

Please do not circulate

Abstract

While Hall (1978) argues that the solution for consumption of a life-cycle model with uncertainty and isoelastic preferences evolves as a random walk in first order approximation around small income shocks, I show that it does not. First, I explain that the solution for consumption departs from a random walk because of precautionary behavior. A prudent household facing uncertainty allocates more of its resources to the uncertain future, raising its future consumption growth to an extent that depends on its current variables. This precautionary component of consumption growth does not disappear in first order approximation around small realized income shocks because it relates to the magnitude of the uncertainty ex-ante, not to the magnitude of the realized shocks ex-post. Second, I account for the previous literature that finds consumption to evolve approximately as a random walk: I show that there is an element of circular reasoning in the derivation of a random walk expression. Third, I assess the importance of the correlation of assets, total income and transitory income with subsequent consumption growth in a calibrated life-cycle model. I find all correlations to be highly significant.

*Economics Department, European University Institute, Via dei Roccettini 9, 50014 San Domenico di Fiesole, Italy; jeanne.commault@gmail.com.

'Indeed, when Hall first presented the paper deriving and testing the random-walk result, one prominent macroeconomist told him that he must have been on drugs when he wrote the paper.'
David Romer, Advanced Macroeconomics, 4th Edition, April 2011, p375

1 Introduction

How does consumption evolve over time when it is chosen as the solution of a life-cycle maximization problem? In a seminal paper published in the *Journal of Political Economy* in 1978, Hall argues that, in a standard life-cycle model, the marginal utility of consumption evolves as a random walk (e.g. its future growth does not depend on any currently observed variables), and consumption itself inherits this characteristic when preferences are quadratic or, in first order approximation around small shocks, when preferences are isoelastic. The random walk property constitutes a convenient restriction that has been and is still abundantly exploited in the literature. Although a number of papers now allow some households to depart from a random-walk behavior because of myopia or liquidity constraints, they still accept that Hall's result applies to some part of the population. More precisely, three streams of the recent literature actively rely on the random walk expression of consumption: i) papers testing the validity of the life-cycle model or measuring the respective shares of random-walk versus constrained consumers in the population;[1] ii) papers relying on the random walk expression or on the first or second order log-linearized Euler equation to interpret numerical simulations of life-cycle models;[2] iii) papers estimating elasticities to income shocks: the random walk expression implies that the log-consumption growth of a household does not depend on the past income shocks it has experienced, which can be used as an identifying restriction to measure separately the impact of transitory and permanent shocks on log-consumption; the method builds on the seminal work of Blundell, Pistaferri, and Preston (2008) (BPP afterwards) and is now influential in multiple fields including household finance, labor, development, and housing.[3]

In this paper, I show that these uses of the random walk restriction need to be reexamined because the consumption solution of a life-cycle model with isoelastic preferences and income uncertainty does not evolve as a random walk, neither exactly nor in first order approximation around small income shocks. Consumption departs from a random walk because the precautionary component of future consumption growth correlates with currently observed variables, in particular with the level of assets that a household

[1] A vast literature attempts to test the life-cycle model by checking whether the random walk restriction holds in consumption data, with mixed results generally attributed to the low power of the test (see Browning and Lusardi (1996) for a review). Campbell and Mankiw (1989) and Zeldes (1989) generalize the model by allowing for hand-to-mouth consumers in the population and estimate their proportion from the importance of departures from a random walk in consumption data, where they use the response of consumption growth to past income growth and past consumption growth as a departure from a random walk; Following their example, a number of macro and micro studies allow for these two types of households and assess their importance in the population: for instance Iacoviello (2004) measures the proportion of hand-to-mouth implied by the response of consumption growth to changes in home value; Card, Chetty, and Weber (2007) measure the proportion implied by the response of job search behavior to an extension in the duration of unemployment benefits; Attanasio and Borella (2014) test the random walk restriction separately on households with high and low educated heads, who might face more financial constraints; Guvenen and Smith (2014) allow consumers to face a borrowing constraint more stringent than the one implied by a simple interdiction to die in debt, and estimate its level from income and consumption data in an approximated life-cycle model.

[2] The buffer-stock interpretation of the life-cycle model with prudent and impatient consumers proposed by Carroll (Carroll (1992), Carroll (1997), Carroll (2001)) is based on the second order version of the log-linearized Euler equation. The buffer-stock description of the dynamics between consumption growth and income growth in such a model is commonly used to interpret numerical simulations, starting with the seminal paper of Gourinchas and Parker (2002).

[3] I survey the diverse uses of this method in Appendix A.

owns. Previous random walk approximations of consumption neglect the precautionary component of consumption growth because the way they derive an expression of consumption is based on circular reasoning: the approximation point they choose already implies some assumptions on the distribution of future consumption that is the variable whose properties they aim to study. I calibrate and simulate a life-cycle model that mimics the US economy to measure the importance of the correlations between consumption growth and currently observed variables. I find all correlations to be significant at 0.01% both at the household level and after aggregating variables by year.

The model is a life-cycle framework in which a finite-lived household maximizes its intertemporal expected utility, subject to a budget constraint and a terminal condition on wealth so that it cannot die in debt. Its utility is isoelastic (or to put it another way, it is in the class of Constant Relative Risk Aversion utility functions). The household has access to a riskless asset to save and borrow, and receives a stochastic income at each period.

The household is uncertain about its future consumption because it is uncertain about its future income. When preferences are isoelastic, marginal utility is convex so the uncertainty of the household about its future consumption raises its expected marginal utility of future consumption. It leads the household to allocate relatively more resources to future consumption, whose expected marginal utility is increased, thus to choose a steeper expected consumption growth: there is a precautionary component to consumption growth. The magnitude of the precautionary consumption growth that a household chooses depends on its expected distribution of future consumption, which is determined by its current assets and income. These variables are therefore of value in predicting subsequent consumption growth. In particular, everything else being equal, the future consumption of a household that enters the current period with more assets is distributed around a higher level of expected consumption, and is more compressed around its expected value. This means that it is distributed around a region where marginal utility is less convex and that it has a smaller variance, two effects that reduce the need for precautionary consumption growth. Thus, entering the current period with more assets is strictly negatively correlated with subsequent consumption growth.

Approximating consumption around the point where the realized shocks to future income are small does not make the precautionary component of consumption growth disappear. It remains in the approximated expression because the household decides to move more resources to the future before the shocks are realized, independently of whether they eventually happen to be large or not. An alternative approximation around the point where the expected variance of future income is zero, so that both expected and realized income shocks are small, does not yield a random walk expression either: it does make the precautionary component of consumption growth disappear, but the first order term of such an approximation is the effect of a small change in the variance of future income on consumption growth, which depends on the strength of precautionary behavior thus correlates with variables that are currently observed.

Existing random walk expressions of consumption neglect the precautionary correlation between future consumption growth and currently observed variables because they are based on an approximation around the point where the marginal utility of future consumption coincides with the marginal utility of current consumption. By selecting this approximation point, one is already making the assumption that future consumption be close to current consumption up to a noise that is mean-zero when the Euler equation holds, that is, making the assumption that there is no precautionary component to consumption growth and consumption evolves as a random walk. The derivation of the random walk expression is therefore based on a circular reasoning because it begins with what it aims to establish. The same issue affects the derivation of a log-linearized Euler equation at all orders of approximation, which is also obtained as an approximation around the point where the marginal utility of future consumption equals the marginal utility of current consumption. Similarly, the random walk expression of log-consumption that is used as an estimating equation in the BPP method is derived by considering an approximation point that is endogenous to future consumption.

To examine whether the departure from a random walk is empirically important, I simulate a life-cycle model that is calibrated to mimic the behavior of U.S. households with a head aged 30-65 between 1978 and 1992. I find that, although consumers are all solving a standard life-cycle model, the correlation of their past assets and past income with their consumption growth is significant at 0.1%. Aggregating variables by year, the correlation of past aggregate assets and past aggregate income with aggregate consumption growth in the economy is significant at 0.1% as well. These two variables explain more than 90% of the variations in aggregate consumption growth over time.

2 Properties of consumption in a life-cycle model

2.1 The baseline life-cycle model

I consider the standard life-cycle model of consumption under uncertainty. A household i maximizes its total expected utility over time, subject to a budget constraint at each period and a terminal condition on wealth:

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

$$s.t. \quad a_{k+1} = (1+r)a_k - c_k + y_k \qquad \forall t \leq k \leq T, \tag{2.2}$$

$$a_T \geq 0 \tag{2.3}$$

The household 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 consumption expenditures c_t. The period utility $u(.)$ is isoelastic: its functional form is $u(c) = \frac{c^{1-\rho}-1}{1-\rho}$. This means that the marginal utility function is strictly positive, decreasing, and convex ($u'(.) > 0$, $u''(.) < 0$, and $u'''(.) > 0$). Future utility is discounted at a factor β. The budget constraint (2.2) states that the household has access only to one type of asset a_t that delivers a risk-free interest rate r to store its wealth from the end of period $t-1$ to the beginning of period t, and it cannot default on its debt during its lifetime. The terminal condition on wealth (2.3) states that the household cannot die with a strictly positive amount of debt. The term y_t denotes the earned income of period t that is stochastic.

2.2 Consumption growth

The first order condition of this maximization problem, known as the Euler equation, is as follows:

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

It states that an optimizing household seeks to equalize its expected marginal utility over time. If it expects one additional unit of consumption to be more valuable at one period than at another, it is optimal to adjust its allocation of consumption until consuming one additional unit at any given period is no longer more valuable than it would be at other periods.[4] The marginal utility of current consumption is weighted by the deterministic amount $\lambda^{-1} = (\beta(1+r))^{-1}$ because of intertemporal substitution motives.[5] The effect of intertemporal substitution can equivalently be expressed as a weight $\lambda^{1/\rho}$ on current consumption: $u'(c)\lambda^{-1} = c^{-\rho}\lambda^{-1} = (c\lambda^{1/\rho})^{-\rho}$. The Euler equation then rewrites: $u'(c_t\lambda^{1/\rho}) = E_t[u'(c_{t+1})]$.

When preferences are isoelastic, marginal utility is strictly convex so the expected marginal utility of consumption is strictly larger than the marginal utility of expected consumption:

$$E_t[u'(c_{t+1})] > u'(E_t[c_{t+1}]).$$

This is because the convexity means that negative and positive shocks to consumption have asymmetric effects on the marginal utility of consumption: a negative shock raises the value of an additional unit of future consumption more than a positive shock of the same magnitude reduces the value of an additional unit of future consumption. On average across the possible states of the world, the eventuality of negative shocks has a stronger effect than the eventuality of positive shocks so the presence of mean-zero shocks raises the expected marginal utility of consumption above the marginal utility of expected consumption.

I combine the inequality with the Euler equation and apply $(u')^{-1}(.)$ to each side:

$$u'(c_t\lambda^{1/\rho}) > u'(E_t[c_{t+1}])$$
$$c_t\lambda^{1/\rho} < E_t[c_{t+1}]$$

The household chooses a level of current consumption $c_t\lambda^{1/\rho}$ that is below its expected future consumption. Indeed, because current consumption is certain while future consumption is uncertain, future consumption benefits from an increased expected marginal utility. Since marginal utility is decreasing in consumption, it leads the household to reduce its current consumption and increase its future consumption to equalize the marginal utility of current consumption with its expected marginal utility of future consumption. Thus, being uncertain about its future consumption induces the household to raise its expected consumption growth. I denote φ_t this additional consumption growth:

$$E_t[c_{t+1}] = c_t\lambda^{1/\rho} + \underbrace{\varphi_t}_{>0}$$

[4] The multiplier on the natural borrowing limit does not enter this expression because the constraint never binds: the household would never put itself in the situation of possibly consuming zero in the future. Indeed, this would be associated with an infinitely large expected marginal utility of consumption, inducing the household to reduce its current consumption until it falls below the natural borrowing limit.

[5] When the discount factor β of the household and the interest factor $(1+r)$ that rewards its postponing consumption do not compensate each other, the household can find it more enjoyable to consume more right away ($\beta(1+r) < 1$), or prefer to postpone part of its current consumption to take advantage of the interest rate and to consume more later ($\beta(1+r) > 1$)

Figure 1: Equalization of expected marginal utility

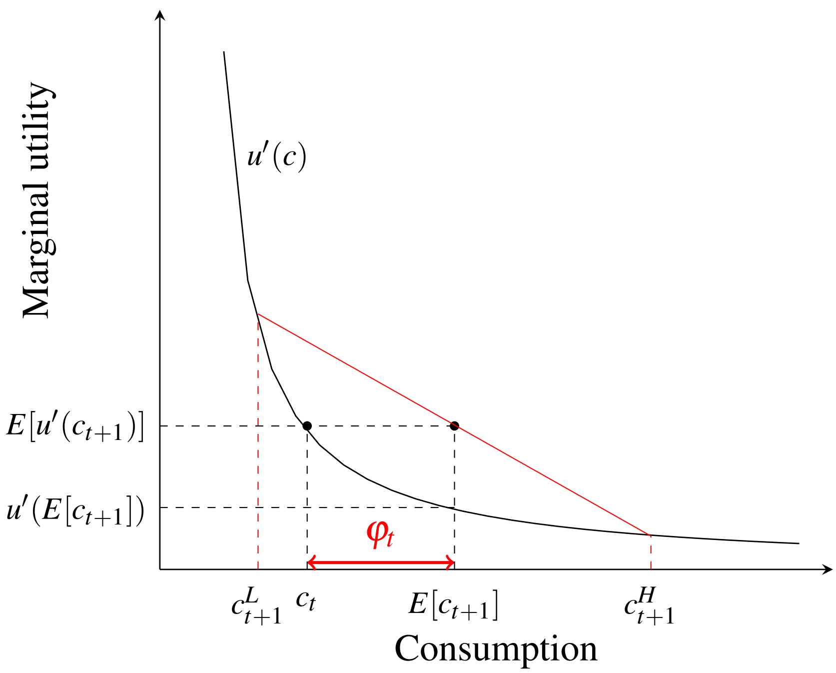

Figure 1. represents how a household who behaves according to a life-cycle maximization problem sets its current and expected consumption. The utility function is isoelastic with $\rho = 2$.

I refer to it as precautionary consumption growth because it corresponds to the effect of uncertainty on consumption growth.[6] Incidentally, precautionary consumption growth coincides with the risk premium of the marginal utility function applied to consumption risk between t and $t+1$.[7]

Figures 1 present a graphical illustration of the mechanism. It pictures a situation in which future consumption can take two values, a low value denoted c^L_{t+1}, and a high value, denoted c^H_{t+1}. Because marginal utility is convex, the increase in marginal utility resulting from a low income realization is much larger than the decrease in marginal utility associated with a high income realization. Taking the average of the two, the expected marginal utility of future consumption is greater than the marginal utility of expected future consumption: $E_t[u'(c_{t+1})]$ is above $u'(E_t[c_{t+1}])$ on the graph. The value c_t is the amount of certain consumption that has the same marginal utility as $E_t[u'(c_{t+1})]$. Since $E_t[u'(c_{t+1})]$ is above $u'(E_t[c_{t+1}])$ and marginal utility is decreasing in consumption, c_t must below $E_t[c_{t+1}]$ and $\varphi_t > 0$: a household facing no uncertainty must consume less than one facing uncertainty on average to have the same expected marginal utility. As a result, the presence of uncertainty makes the expected consumption path steeper.

[6]This is true independently of how the world without uncertainty is defined. Whether income is certain and equals to its expected value or certain and equal to another value, a world without uncertainty implies that the household will seek to equalize $c^{perf.for.}_{t+1} = c^{perf.for.}_t \lambda^{1/\rho}$, thus a zero consumption growth (up to the trend $\lambda^{1/\rho}$.

[7]Kimball (1990) is the first to refer to the risk-premium of the marginal utility as the precautionary premium. He does not make the link between this premium and the value of precautionary growth, but shows that the premium coincide with the additional wealth that a household facing income uncertainty would need to have to consume as much as he would under perfect foresight.

Eventually, consumption evolves as follows:

$$c_{t+1} = c_t \underbrace{\lambda^{1/\rho}}_{\substack{\text{inter.}\\ \text{substitution}}} + \underbrace{\varphi_t}_{\text{precaution}} + \underbrace{\xi_{t+1}}_{\text{innovation}}$$

with $\xi_{t+1} = c_{t+1} - E_t[c_{t+1}]$ the innovation to consumption between t and $t+1$. It shows that future consumption is the sum of current consumption weighted by $\lambda^{1/\rho}$ to account for a household's impatience or willingness to take advantage of high interest rates, plus the precautionary consumption growth that compensates for the impact of bad shocks whose possibility raises expected marginal utility, and an innovation term that captures the update on its consumption that a household makes after receiving income shocks at $t+1$.

3 Departure from a random walk

3.1 Variations of precautionary consumption growth with variables observed at t

Precautionary consumption growth depends both on the functional form of marginal utility and on the distribution of future consumption. A number of variables observed at t influence the distribution of future consumption so they should in general affect the value of precautionary consumption growth, inducing a departure from a random walk.

Assets I prove that for a general class of income processes the precautionary consumption growth of a household decreases strictly with its level of assets.

Theorem: In the model presented above, when income has a flexible transitory-permanent specification ($y_t = y(p_t, \varepsilon_t)$, with $p_t = p(p_{t-1}, \eta_t)$ where ε_t and η_t are the transitory and permanent shocks and $y(.)$ and $p(.)$ are flexible functions) the precautionary consumption growth, φ_t, is negatively correlated with net assets. At any period $0 < t < T$, and for any $0 < k < t$:

$$\frac{d\varphi_t}{da_t} < 0.$$

Intuition: The exact proof is presented in Appendix B. Here, I detail the intuition of the proof with an approximation. I apply the local approximation of Arrow (1965) and Pratt (1964) to the precautionary premium associated with consumption risk at $t+1$ to obtain the following decomposition:

$$\varphi_t = \frac{1}{2} \underbrace{\frac{u'''(E_t[c_{t+1}])}{-u''(E_t[c_{t+1}])}}_{\text{absolute prudence}} \times \underbrace{var_t(c_{t+1})}_{\text{consumption risk}} + o(var_t(c_{t+1})).$$

The first term, the coefficient of absolute prudence, measures the local convexity of the marginal utility function at the point where $c = E_t[c_{t+1}]$. Therefore, it captures the extent to which fluctuations in consumption translates into fluctuations in marginal utility around $E_t[c_{t+1}]$, and the fluctuations in marginal utility are eventually what raise the household's expected marginal utility of future consumption inducing it to move resources from the present to the future. The second term, the variance of future consumption, measures the magnitude of the fluctuations in future consumption thus the risk towards its future consumption that the household faces. I take the derivative with respect to net assets of this approximation of φ_t:

$$\frac{d\varphi_t}{da_t} \approx \underbrace{\overbrace{\frac{dE_t[c_{t+1}]}{da_t}}^{>0} \overbrace{\left(\frac{u'''}{-u''}\right)'(E_t[c_{t+1}])}^{<0}}_{\text{change in prudence}} var_t(c_{t+1}) + p(E_t[c_{t+1}]) \underbrace{\overbrace{\frac{dvar_t(c_{t+1})}{da_t}}^{<0}}_{\text{change in risk}} < 0.$$

The impact of a gain in net assets decomposes into two effects. First, it decreases the value of absolute prudence. This is because a gain in net assets increases future expected consumption $E_t[c_{t+1}]$, and the coefficient of absolute prudence is decreasing in consumption (a property of isoelastic preferences): around higher levels of consumption, the same fluctuations in consumption induce less fluctuations in marginal utility. Second, a gain in net assets decreases the variance of future consumption. Intuitively, such a gain does not only shift the distribution of consumption around a higher expected value, it also compresses its distribution. Indeed, when preferences are isoelastic and income has a transitory-permanent specification, a gain in assets does not affect consumption identically in all states of the world but raises more consumption when adverse transitory or permanent shocks realize, in the low consumption states of the world, and raises it less when favorable transitory or permanent shocks realize, in the high consumption states of the world, reducing the variance of c_{t+1}. Thus, because a gain in net assets diminishes both the extent to which fluctuations in consumption translate into fluctuations in marginal utility (measured by the local value of absolute prudence) and the level of consumption fluctuations (measured by the variance of future consumption), they imply a decrease in precautionary saving.

Figures 2(a) and 2(b) illustrate this mechanism graphically. The top figure 2(a) represents the first effect of a gain assets: a shift upwards in the distribution of future consumption, keeping the distance between the low and high values of future consumption constant. It induces a decrease in the amount of precautionary consumption growth necessary to compensate for the presence of uncertainty towards c_{t+1}, because around a higher level of consumption the convexity of marginal utility is less pronounced: the convexity of marginal utility below the red distribution is more pronounced than below the blue distribution. Thus, following a shift upwards, the same possible variations in future consumption do not raise the expected marginal utility of future consumption as much as before and a smaller amount of precautionary consumption growth is sufficient to equalize the current and expected future marginal utility of consumption. The bottom figure (b) accounts for the second effect of a gain in assets: a compression in the distribution of future consumption. It induces an additional decrease in the need for precautionary saving, because the variations in future consumption are smaller: the orange distribution lies within the blue distribution.

In the general case when I make no specific assumptions on the specification of income, I do not prove that the precautionary component of consumption growth is always negatively correlated with the level of assets of the household, but I show that it correlates strictly with the level of assets of the household, except for a knife-edge parametrization and history of past shocks such that the change in consumption risk caused by a gain in assets compensates exactly the change in absolute prudence.

Income A change in current income modifies the distribution of future consumption as well, but does not necessarily result in a shift upwards combined with a compression of the distribution around the expected

Figure 2: Equalization of expected marginal utility before and after a gain in assets

(a) Effect 1 - Shift up in the distribution of future consumption

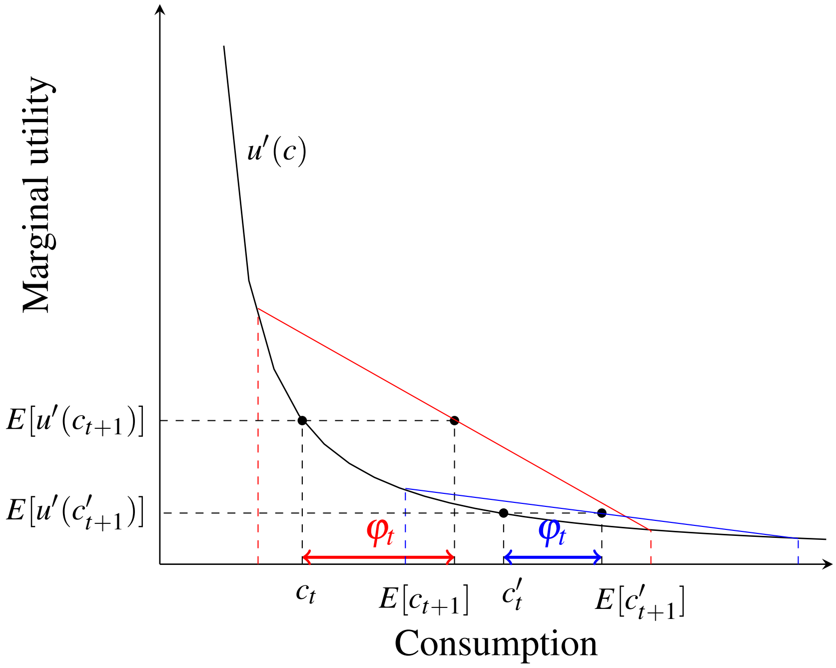

(b) Effect 2 - Compression in the distribution of future consumption

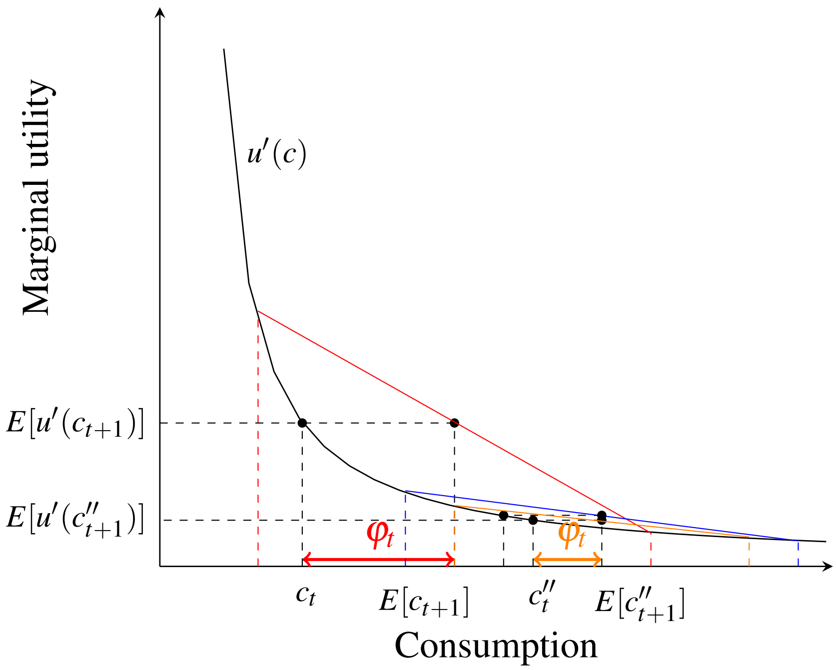

Figure 1. represents how a household who behaves according to a life-cycle maximization problem changes its current and expected consumption following an exogenous gain in assets. The utility function is isoelastic with $\rho = 2$. The change decomposes into two effects: i) a gain in assets shifts the distribution of future consumption up (top figure); ii) a gain in assets compresses the distribution around its expected value (bottom figure).

value. In the case when it does, such as when the shock is an MA(0) that increases current income and does not affect the value of future income, then it triggers a decrease in precautionary consumption growth for the same reasons as a gain in assets. In general, an income shock modifies the distribution of future consumption but in a way that can either reduce or increase precautionary consumption growth so the direction of the effect is undetermined.

3.2 Departure of consumption from a random walk with trend

What is a random walk process? A variable c evolves as a random walk with trend α if the quasi-difference $(c_{t+1} - \alpha c_t)$ is independent of the current and past values of c, thus independent of any variable x_{t-s} ($s \geq 0$) that correlates with the current and past values of c. The independence implies the non-correlation: the expected product of quasi-consumption growth with current or past variables is the product of their expected values. Observing x_{t-s} has no effect on one's expected value of $(c_{t+1} - \alpha c_t)$:

$$c \text{ is a random walk with trend } \alpha \;\Rightarrow\; E[(c_{t+1} - \alpha c_t)x_{t-s}] = E[c_{t+1} - \alpha c_t]E[x_{t-s}]$$

with $E[.]$ the unconditional expectation operator, which corresponds to the sum over the possible realizations of the past shocks that the household has received between period 0 and period $t+1$. It means that in a population of households making their decisions independently and receiving shocks drawn independently, if c evolves as a random walk with trend α, the values of $(c_{t+1} - \alpha c_t)$ and of x_{t-s} are uncorrelated. Thus, a regression of $(c_{t+1} - \alpha c_t)$ over a variable x_{t-s} in this population should yield a zero coefficient (or equivalently a regression of c_{t+1} over c_t and x_{t-s} should yield a zero coefficient on x_{t-s}).[8] This is the test that Hall (1978) uses to examine whether consumption evolves as a random walk.

Now, is the solution of the model such that $(c_{t+1} - \lambda^{1/\rho} c_t)$ does not correlate with any variable x_{t-s}? The consumption growth of a household adjusted for a trend $\lambda^{1/\rho}$ is the sum of its precautionary consumption growth plus an innovation of mean zero:

$$c_{t+1} - c_t \lambda^{1/\rho} = \underbrace{\varphi_t}_{\substack{\text{correlates with} \\ \text{variables at } t}} + \xi_{t+1}$$

I examine the unconditional expected product of the quasi-difference in consumption of a household and its level of assets:

$$E[c_{t+1} - c_t \lambda^{1/\rho} a_t] = E[\varphi_t]E[a_t] + \underbrace{cov(\varphi_t, a_t)}_{\neq 0} + E[\xi_{t+1}]E[a_t]$$

When income has a flexible transitory-permanent specification, the realizations of past and current transitory shocks influence both the amount of assets that a household owns and its precautionary consumption growth (through their effect on assets and income, which modifies the distribution of future consumption). The covariance between precautionary consumption growth and assets is non-zero (if the covariance was zero for any initial level of assets, for a given distribution of income it would be true for any distribution fo income, nana, avec un niveau d'asset initial diffrent, at augmente, varphi diminue, et la covariance bouge du moins je pense). It is not necessarily negative, however, because although the

[8] Because the estimator of the coefficient is $\hat{\beta}_{x_t} = \frac{E[(c_{i,t+1} - \alpha c_{i,t})x_{i,t}] - E[c_{i,t+1} - \alpha c_{i,t}]E[x_{i,t}]}{E[x_{i,t}^2] - E[x_{i,t}]^2} = 0$.

realizations of past transitory shocks correlate positively with current assets and negatively with precautionary consumption growth, inducing a negative covariance between the two, the realizations of past permanent shocks can raise current assets and raise also precautionary consumption growth if they make future income sufficiently more risky, possibly inducing a positive covariance between the two. In the general case when I make no specific assumptions on the specification of income, the covariance is non-zero in general, except at particular level of assets and for particular sets of shocks such that at this point the impact of shocks that raise assets and decrease precautionary consumption growth exactly compensate the impact of shocks that raise assets and decrease precautionary consumption growth. Thus, the expected product of trend-adjusted consumption growth and current assets is not the product of their expected values and the trend-adjusted consumption growth correlates with its level of assets are not independent.

In addition, although the innovation ξ_{t+1} does not correlate with the level of assets and income that a household owns, they are not in general independent. The innovation measures the revision in consumption that a household makes at $t+1$ upon observing the income shocks it receives. Its expected value is always zero so it does not vary with any current or past variable x_{t-s} (there is non-correlation). Yet, the innovation has a precautionary component: the revision in precautionary saving that the household makes upon observing the shocks it receives. The distribution of this precautionary component around its expected value of zero depends on the distribution of future consumption thus varies with current assets and income in general (there is no independence). I show it formally in Appendix C. Thus, even the result that marginal utility evolves as a random walk does not hold in a life-cycle model with uncertainty and isoleastic preferences.

3.3 Departure of approximated consumption from a random walk

Approximation around small income shocks The expression resulting from a first-order approximation of consumption at the point where income equals its expected value does not describe a random walk either. I take a first order approximation at the point where $y_{t+1} = E_t[y_t]$ of the expression of consumption derived from the first-order condition:

$$c_{t+1} - c_t \lambda^{1/\rho} \approx \underbrace{\varphi_t}_{\substack{\text{correlates with}\\ \text{variables at } t}} + \xi_{t+1|y_{t+1}=E_t[y_{t+1}]} + (y_{t+1} - E_t[y_t]) \frac{d\xi_{t+1}}{dy_{t+1}}_{|y_{t+1}=E_t[y_{t+1}]}$$

The precautionary component of consumption growth φ_t does not disappear. This is because precautionary consumption growth is decided at t and depends on the possibility that income shocks at $t+1$ will shift consumption away from its expected value, but it does not depend on whether such shocks are eventually realized or not. Thus, it is unaffected by an approximation around small realized innovations to income. Because the precautionary component of consumption growth correlates in particular with the level of assets that a household owns, it makes consumption growth partly predictable from variables observed at t.

The same intuition holds more generally for approximations around the expected value of any components of income, and not only for approximations around the expected value of income itself. Without loss of generality, future income can be expressed a function of a vector of shocks ω^{t+1}, $y_{t+1} = g_{t+1}(\omega^{t+1})$. I take a first approximation at the point where each term of the vector equals its expected value at t: $\omega^{t+1} = E_t[\omega^{t+1}]$. The precautionary component of consumption growth φ_t remains, inducing a departure from a random walk because it correlates with variables observed at t.

Approximation around small income variance The reason why the approximations above do not eliminate precautionary behavior is because, although the realized income or the realized shocks are close to their expected values, the household does not know it ex-ante and still anticipates the possibility of large shocks. I consider an approximation around the point where the variance at t of income at $t+1$ is zero, which means that innovations to future income are small and the household anticipates that they will be so. I denote $\sigma_t^{y_{t+1}} = 0$ the variance at t of income at $t+1$. The approximated expression of consumption growth is:

$$c_{t+1} - c_t \lambda^{1/\rho} \approx (\sigma_t^{y_{t+1}} - 0)(\underbrace{\frac{d\varphi_t}{d\sigma_t^{y_{t+1}}}}_{\substack{\text{correlates with} \\ \text{variables at } t}} + \frac{d\xi_{t+1}}{d\sigma_t^{y_{t+1}}})_{|\sigma_t^{y_{t+1}}=0}$$

Indeed, when the variance of future income is zero, expected income and actual income coincide, and so do expected consumption and actual consumption. Because there are no shocks moving consumption above or below its expected value, there is no precautionary consumption growth between t and $t+1$. Thus, at the approximation point, both the precautionary component of consumption growth and the innovation are zero. However, the first order term of the approximation correlates in general with variables observed at t. It depends on the response of precautionary consumption growth to a change in the variance of future income: if it did not then precautionary consumption growth would not depend on the level of assets, at least for small levels of variance, which is not the case. This approximation of consumption growth does not yield a random walk process either.

4 Circular reasoning in expressions of consumption growth

4.1 Hall's (1978) random walk expression

In his seminal 1978 paper, Hall considers the same life-cycle model as introduced here. Hall's expression is a first order approximation of future consumption $c_{t+1} = (u')^{-1}(u'(c_{t+1}))$ around the point where the marginal utility of future consumption coincides with its expected value, that is, with the marginal utility of current consumption $u'(c_{t+1}) = E_t[u'(c_{t+1})] = u'(c_t\lambda^{1/\rho})$:

$$c_{t+1} = ((u')^{-1}(u'(c_{t+1}))$$
$$c_{t+1} \approx (u')^{-1}(u'(c_t\lambda^{1/\rho})) + (u'(c_{t+1}) - u'(c_t\lambda^{1/\rho}))\frac{dc_{t+1}}{du'(c_{t+1})}_{|u'(c_{t+1})=u'(c_t\lambda^{1/\rho})}$$
$$c_{t+1} \approx c_t\lambda^{1/\rho} + \frac{u'(c_{t+1}) - E_t[u'(c_{t+1})]}{u''(c_t\lambda^{1/\rho})} \quad , \quad E_t[u'(c_{t+1}) - E_t[u'(c_{t+1})]] = 0$$

where $u'(c_{t+1}) - E_t[u'(c_{t+1})]$ is the innovation to marginal utility between t and $t+1$. The expression implies that consumption evolves as a random walk with trend $\lambda^{1/\rho}$, because future consumption growth adjusted for trend $\lambda^{1/\rho}$ does not correlate with any current or past variables: for any x_{t-s}, $E[c_{t+1} - c_t\lambda^{1/\rho}|x_{t-s}] = E[u'(c_{t+1}) - E_t[u'(c_{t+1})]|x_{t-s}] = 0 = E[u'(c_{t+1}) - E_t[u'(c_{t+1})]] = E[c_{t+1} - c_t\lambda^{1/\rho}]$.

What is the problem? This expression is obtained by approximating future consumption around the point where the marginal utility of future consumption is close to its current value. By selecting such an approximation point Hall is already making assumptions on future consumption: he is imposing that c_{t+1} is distributed in such a way that $u'(c_{t+1})$ is close to $u'(c_t\lambda^{1/\rho})$ in all states of the world. It means assuming that c_{t+1} is distributed around $c_t\lambda^{1/\rho}$[9], that is, assuming that the precautionary premium φ_t is zero (or that its effect always compensates exactly the innovation, $\xi_{t+1} \approx -\varphi_t$ for all y_{t+1}, which comes down to the same assumption[10]), therefore assuming that consumption is close to a random walk. The resulting expression is based on circular reasoning as the result is implicitly assumed in the premises. In the unconstrained model, the approximation point selected by Hall is never reached, because the precautionary component φ_t is strictly positive in all states of the world.

With this same method it is in fact equally possible to approximate consumption as the sum of a different value $c^* \neq c_t\lambda^{1/\rho}$ plus an innovation of mean zero.[11] Which choice of constant gives a better representation of the actual consumption process is unclear because the quality of the representation is determined by the quality of the approximation point, which depends on how the actual consumption process evolves; that is, what the researcher is trying to determine.

4.2 The log-linearized Euler equation

The log-linearized Euler equation is a popular expression of log-consumption growth, used in both the micro and macro literature:[13]

$$\Delta ln(c_{t+1}) \approx \frac{1}{\rho} ln(\beta(1+r)) + \frac{1}{\rho}\frac{1}{2} var_t\left(\frac{\varepsilon_{t+1}}{u'(c_t)}\right) + \mu_{t+1}$$

with $\mu_{t+1} = \Delta ln(c_{t+1}) - E_t[\Delta ln(c_{t+1})]$ the innovation to log-consumption growth. It implies that log-consumption evolves as a random walk (if the variance term is neglected or assumed to be constant), or as a random walk plus the time-varying variance term. The literature review of Browning and Lusardi (1996) presents the common way to derive this expression. It is obtained as the second order approximation of the log of future consumption around the point where the marginal utility of future consumption equals the marginal utility of current consumption: $u'(c_{t+1}) = u'(c_t\lambda^{1/\rho})$. It is therefore subject to the

[9]Future consumption c_{t+1} is implicitly assumed to approximately write as the sum of $c_t\lambda^{1/\rho}$ plus a noise of mean zero because the difference between c_{t+1} and $c_t\lambda^{1/\rho}$ is implicitly assumed to be approximately a scalar of the difference between $u'(c_{t+1})$ and $u'(c_t\lambda^{1/\rho})$, which is of mean-zero when the Euler equation holds.

[10]Indeed, since the innovation is independent from variables observed at t so after applying the expectation operator it implies $0 = E_t[\xi_{t+1}] = -E_t[\varphi_t] = -\varphi_t$.

[11] For any arbitrary function $f(.)$ that is continuous over the distribution of c_{t+1}, there exists a c^* such that $E_t[f(c_{t+1})] = f(c^*)$. The value of c^* is strictly smaller than $E_t[c_{t+1}]$ is $f(.)$ is convex, strictly larger than $E_t[c_{t+1}]$ if $f(.)$ is concave and equal to $E_t[c_{t+1}]$ if $f(.)$ is linear.[12] If $c_t\lambda^{1/\rho}$ is smaller than $E_t[c_{t+1}]$, one can pick a $f(.)$ that is concave to be sure that $c^* \neq c_t\lambda^{1/\rho}$, or conversely pick a convex $f(.)$ if $c_t\lambda^{1/\rho}$ is larger than $E_t[c_{t+1}]$. I select a function $f(.)$ such that the associated c^* is different from $c_t\lambda^{1/\rho}$ and take a first order approximation of the trivial equilibrium relationship $c_{t+1} = f^{-1}(f(c_{t+1}))$ around the point where $f(c_{t+1}) = f(c^*)$:

$$c_{t+1} = (f)^{-1}(f(c_{t+1}))$$

$$c_{t+1} \approx (f)^{-1}(f(c^*)) + (f(c_{t+1}) - f(c^*))\frac{dc_{t+1}}{df(c_{t+1})}$$

$$c_{t+1} \approx c^* + (f(c_{t+1}) - E_t[f(c_{t+1})])\frac{1}{f'(c^*)} \quad , \quad E_t[f(c_{t+1}) - E_t[f(c_{t+1})]] = 0$$

prendre deux fonctions u1 et u2 avec u2 plus convexe que u1

[13]Most versions additionally includes deterministic demographic changes. To simplify the presentation, I do not include them but their presence does not alter the point that I making.

same problem as the approximation of Hall (1978): it is based on the implicit assumption that future log-consumption is distributed around current log-consumption plus a noise of mean zero, neglecting the precautionary component of consumption growth. Although the approximation point is the same as in Hall, the assumption on the distribution of future consumption is slightly different because the approximation is applied to log-consumption and not consumption.

Tests of whether the data show support for the random walk with drift version of the log-linearized Euler equation show mixed results. The blame for such mixed results is commonly put on the fact that the random walk with drift is neglecting the time-varying variance term (Carroll (1992), Carroll (1997), Carroll (2001)). As a solution, a number of papers including those of Carroll favor the use of the second order version with time-varying variance for testing the model, estimating its parameters, or interpreting the results from numerical simulations.[14] Yet, the second order log-linearized Euler equation with time-varying variance relies on the same circular reasoning as its first order version. A life-cycle model does not imply that log-consumption growth resembles its second order log-linearized expression, nor that the contribution of precautionary behavior coincides with the variance term $\frac{1}{2}var_t\left(\frac{\varepsilon_{t+1}}{u'(c_t)}\right)$. The issue with the log-linearized Euler equation is not that some higher-order terms are neglected but that it is based on arbitrary foundations, which cannot be corrected by taking into account a finite number of additional higher order terms.

4.3 Log-consumption growth and income shocks: the BPP estimating equation

Blundell, Pistaferri, and Preston (2008) design a mechanism to estimate consumption elasticities to income shocks based on a detailed version of the log-linearized Euler equation. They take a precise stand on the specification of income, namely that it follows a transitory-permanent income process with ε and η the transitory and permanent income shocks, and make explicit the impact of future shocks on future log-consumption growth:

$$\Delta ln(c_{t+1}) \approx \frac{1}{\rho}\beta(1+r) + \Phi^{\varepsilon}\varepsilon_{t+1} + \Phi^{\eta}\eta_{t+1}$$

Based on this expression, they devise a method to identify the elasticities of consumption to the transitory and permanent income shocks Φ^{ε} and Φ^{η}.

The derivation of this expression requires implicit assumptions that are similar to those at the root of Hall's expression, plus additional ones. More precisely, Blundell, Pistaferri, and Preston (2008) base the derivation of the expression on Blundell, Low, and Preston (2013), who obtain it by taking three successive approximations: i) around the point where $ln(c_{t+1})$ equals $ln(c_t\lambda_t^{1/\rho})$ ii) around the point where $ln(c_{t+1})$ equals $E_t[ln(c_{t+1})]$ iii) around the point where the realized income shocks are zero $\varepsilon_{t+1}=0$ and $\eta_{t+1}=0$. I detail in Appendix D the steps of their derivation.

Approximation i) is problematic for the same reason as in the derivation of Hall (1978): the assumption that $ln(c_{t+1})$ is close in general to $ln(c_t\lambda_t^{1/\rho})$ implies an assumption on the distribution of future consumption, which is the variable whose behavior is to be examined. In addition, the identity between

[14] Carroll (2001) notes that the second order log-linearized Euler equation comes with issues as well though: he finds that even estimations that accounts for the second order variance term fail to consistently measure some parameters of simulated data and attributes it to the effect of some higher order terms that are being neglected. He still advocates the use of the second order equation, however, in the case of idiosyncratic level data (a discussion that appears in the 'good ideas' section of the paper).

approximations i) and ii) does not hold in the model: it would only do so if the precautionary component of consumption growth φ_t was zero while it is strictly positive for isoelastic preferences. The resulting expression is identical to that of a log-linearized Euler equation although the initial approximation point is not exactly the same ($ln(c_{t+1}) = ln(c_t \lambda_t^{1/\rho})$ instead of $u'(c_{t+1}) = u'(c_t \lambda_t^{1/\rho})$), thanks to this additional assumption that approximations i) and ii) coincide (e.g. that the point where $ln(c_{t+1}) = ln(c_t \lambda_t^{1/\rho})$ is the point where $ln(c_{t+1}) = E_t[ln(c_{t+1})]$). There is not reason why the identity between approximations ii) and iii) should hold either: whenever log-consumption is not a linear function of log-income, the value of expected log-consumption (when the innovation to log-consumption is zero) differs from the value of consumption at expected log-income (when income innovation to log-income is zero).

Thus, this expression results from implicit assumptions that eliminate the terms making log-consumption growth history-dependent and non-linear in the income shocks.

5 Quantitative importance of the departure from a random walk

5.1 Simulation of a life-cycle model

How quantitatively important is this departure from a random walk model? To answer this question, I simulate the life-cycle model described in the first section. The calibration is the same as that of Kaplan and Violante (2010) who seek to reproduce the data used by Blundell, Pistaferri, and Preston (2008), e.g. U.S. households between 1978 and 1992 whose head is aged 30-65. The model period is one year. Households enter the labor market at age 25. They work for $T^w = 35$ years and retire at age 60. Past that age, they have a stochastic probability of dying at each period until age 95 when they die with certainty. The survival rates are obtained from the National Center for Health Statistics (1992). The utility function is isoelastic with parameter $\rho = 2$. Since the model is calibrated to generate only half of the total wealth in the US economy (the PSID and CEX survey data undersample the top of the wealth distribution), the interest rate is not determined in equilibrium but set at r = 3 percent. The discount factor β is such that $\beta = \frac{1}{1+r} \approx 0.98$. The common deterministic age profile for log income is calibrated using PSID data. All households start their economic life with zero wealth, $a_0 = 0$. At each period of their working life, households receive a net income y_t. The three parameters of the income process are the variance of the two shocks, σ_ε and σ_η, and the variance in the population of the initial value of the permanent component σ_{z0}. The variance of transitory shocks is set at 0.05. I kept the value used by Kaplan and Violante (2010) although their argument is that it matches the point estimate of Blundell, Pistaferri, and Preston (2008) although my discussion shows that these estimates are possibly biased: the initial bias concerns the pass-through consumption coefficients, but the variance estimates can be affected as well since the two are jointly estimated. The variance of permanent shocks is set at 0.01 to match the rise in earnings dispersion over the life cycle in the PSID from age 25 to age 60. The initial variance of the permanent shocks is set at 0.15 to match the dispersion of household earnings at age 25. After retirement, households receive a pension that is based on the gross income they have received, itself obtained by applying the inverse of a tax function to net income.[15]

[15]The pension function is designed by Kaplan and Violante to mimic the U.S. system. They specify that benefits are equal to 90 percent of average past gross earnings up to a given bend point, then 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 cross-sectional average gross earnings, based on the US legislation and individual earnings data for 1990. Benefits are then scaled proportionately so that a worker earning average labor income each year is entitled to a replacement rate of 45 percent. Gross earnings are

The model is solved using the method of endogenous grid points developed by Carroll (2012).[16] I simulate an artificial panel of 5,000 households for 70 periods, and keep the observations of the 35 periods of working life. This leaves me with 175,000 household-year observations of consumption and 170,000 household-year observations of consumption growth.

5.2 Regressions on simulated data

Table 1: Descriptive statistics

	a_t	y_t	c_t	Δc_{t+1}	$\Delta^{35} c_{t+1}$
Mean	$51,647$	$31,522$	$27,434$	253	$8,613$
Std. Dev.	$104,647$	$22,457$	$15,492$	$15,694$	$2,380$
Obs.	$175,000$	$175,000$	$175,000$	$170,000$	$5,000$

Note: This table presents the average level of assets, yearly income, yearly transitory income shock, yearly consumption, yearly consumption growth, and life-time consumption growth in the simulated population.

Table 1 presents some statistics on the assets, income, consumption, and consumption growth of the households with a head aged 25-60 of this economy. The average amount of assets that these households owned at the beginning of a year period is $51,647$ dollars, their average income over a year is $31,522$. The average value of the transitory shock they receive is mechanically very small, because the average population value is closed to its expected value that is normalized to zero. Precautionary behavior is the only source of expected consumption growth in this model, because the effect of the discount factor and the interest rate cancel out ($\beta(1+r) = \lambda = 1$). It generates an average yearly consumption growth of 253 dollars and an average consumption growth of $8,613$ over the 35 years period of the working life, which corresponds to a percentage consumption growth of 39.3% over the life-cycle.

Table 2: Effect of past assets and past income on consumption growth

Δc_{t+1}	Household-level	Aggregate
a_t	-0.0017^{***} (0.0001)	-0.0022^{***} (0.0003)
y_t	0.0094^{***} (0.0002)	0.0109^{***} (0.0006)
R^2	0.0134	0.9052
Obs.	$170,000$	34

Note: This table presents the results from a regression of consumption growth on assets, income, and transitory income shock in the simulated data. The first column uses variables at the household-level; the second column uses yearly aggregates of these variables. Standard errors are in parenthesis. $***$ significant at 0.1%.

computed by inverting the nonlinear tax function estimated by Gouveia and Strauss (1994) and applying it to net earnings.

[16] The number of gridpoints for each variable is the same as used in Kaplan and Violante: 100 exponentially spaced grid points for assets, 19 points for lifetime average earnings, 39 equally spaced points for the permanent component of income, and 19 equally spaced points for the transitory component of income.

Table 2 shows the results from a regression of consumption growth on assets, income, and the transitory income shock in the simulated data. In the first column, the variables are at the household-level. Although households behave according to the standard life-cycle model without borrowing constraints, the correlations between consumption growth and past assets or past income are significant at 0.1%. At a given level of income and transitory income shock, entering the year with $\$1,000$ more in assets induces a household to choose a consumption growth that is smaller by \$1.7, which corresponds to a 0.7% decrease in average consumption growth. This is consistent with the theoretical prediction that in this model, everything else being equal, a gain in assets triggers a decrease in precautionary consumption growth.[17] At a given level of assets and transitory income shock, earning $\$1,000$ more over the year induces a household to choose a consumption growth that is larger by \$9.4, which corresponds to a 3.7% increase in average consumption growth. A change in current income can be driven either by a change in transitory income, that would reduce the need for precautionary saving, or in permanent income, whose effect on precautionary saving is undetermined. This result suggests that, in this model and for this calibration, the effect on the need for precautionary growth of the increased income risk associated with a gain in permanent income dominates the effect of the increased income level it entails: a gain in permanent income strengthens the need for precautionary growth, so it raises consumption growth. The joint observation of assets, income and transitory income shocks explain 1.4% of the variance of consumption growth in the population, which means that most of the household-level variations in consumption growth between t and $t+1$ is driven by the shocks that households receive at $t+1$.

The second column presents the results obtained with aggregated variables. At the aggregate level as well, the correlation between consumption growth and past assets or past income are significant at 0.1%. The point estimates are little larger. At a given level of income and transitory income shock, an economy in which all households enter the year with $\$1,000$ is associated with a \$2.2 smaller aggregate consumption growth, which corresponds to a 0.9% decrease. At a given level of assets and transitory income shock, an economy in which all households earn a $\$1,000$ more over the year is associated with a \$10.9 larger aggregate consumption growth, which corresponds to a 4.3% increase. The joint observation of aggregate assets and income explains 90.5% of the variance of aggregate consumption growth across periods, which is much larger than at the household-level. Indeed, the innovation component of consumption growth is now mostly averaged away in the population, so the observed variations in aggregate consumption growth reflect variations in expected consumption growth. Because in this model expected consumption growth is entirely caused by precautionary behavior, the determinants of precautionary behavior that are assets, current income, and transitory income explain most of the variations in aggregate consumption growth.

6 Conclusion

Consumption does not evolve as a random walk when modeled as the solution of an intertemporal maximization problem with income uncertainty. This is because in this model there is a precautionary component to the consumption growth of a household that depends on its current variables in particular on the amount of assets that it owns and on its income. The reason why future consumption appears

[17] Because with a transitory-permanent income process current income and transitory income determine entirely permanent income thus the distribution of future income, controlling for current income and transitory income is equivalent to having everything else being equal.

to be non-correlated with current and past variables (up to a multiplicative trend) in previous studies is because those studies are based on an endogenous approximation of consumption at the point where future consumption equals its current value.

This has implications in particular for the estimation of the proportion of households that are credit constrained (commonly measured as the proportion for whom consumption does not obey a random walk), for the interpretation of numerical simulation results, and for the estimations of household-level elasticities related to consumption and saving behavior. I explore in Commault (2018) the consequence of the departure from a random walk for the estimators of the elasticity of consumption to transitory and permanent income shocks, and develop a robust estimator of the elasticity of consumption to transitory income shocks.

Appendix A The BPP estimator in the literature

In household finance studies, Kaufmann and Pistaferri (2009) generalize the BPP method to account for advance information of consumers; Casado (2011) implements the BPP estimator in a database of Spanish households; Blundell, Low, and Preston (2013) adapt it to the use of cross-sectional data and to a more general income process; Hryshko (2014) allows for a correlation between the transitory and permanent shocks; Etheridge (2015) uses the BPP estimator to disentangle between rival specifications of income; Bayer and Juessen (2015) apply it to estimate the response of happiness to transitory and permanent income shocks; Gosh (2016) generalizes the income process and extends the BPP method to exploit third moments together with second moments of income and consumption growth.

In labor, Ortigueira and Siassi (2013) and Heathcote, Storesletten, and Violante (2014) use the BPP estimates as a benchmark to which they compare their simulation results; Blundell, Pistaferri, and Saporta-Eksten (2016) allow for endogenous labor supply and estimate its elasticity to transitory and permanent shocks on a worker's own wage rate and on that of the worker's spouse; Pistaferri, Saporta-Eksten, and Blundell (2017) apply the method to estimate the elasticity of hours spent with children as well.

In development, Zheng and Santaeulalia (2016) estimate the evolution of the elasticity of consumption to income shocks during the period of large and sustained GDP growth in China, Attanasio, Meghir, and Mommaerts (2015) compare the elasticity of consumption to shocks at the village level and at the individual level to assess the importance of insurance mechanisms within the village.

In housing, Carlos Hatchondo, Martinez, and Sànchez (2015) compare the consumption elasticity simulated from a model with mortgage default to the BPP estimates, and Hedlund, Karahan, Mitman, and Ozkan (2017) use the BPP method to measure the elasticity of consumption of households with different leverage ratios.

Appendix B Proof of the Theorem

I prove that at any period $0 < t < T$, and for any $0 < k < t$:

$$\frac{d\varphi_t}{da_t} < 0.$$

Because precautionary consumption growth is the difference between current and expected future consumption, precautionary consumption growth decreases with assets if it is optimal for a household to increase more its current consumption than its expected future consumption in response to a gain in assets:

$$\frac{dE_t[c_{t+1}]}{da_t} = \frac{dc_t\lambda^{1/\rho}}{da_t} + \frac{d\varphi_t}{da_t}$$

To compare the response of current and expected future consumption, I derive each side of the Euler equation with respect to a_t and divide by $(-u''(c_t\lambda^{1/\rho}))$:

$$\frac{dc_t\lambda^{1/\rho}}{da_t} = E_t[\frac{dc_{t+1}}{da_t}\frac{-u''(c_{t+1})}{-u''(c_t\lambda^{1/\rho})}]$$

$$\frac{dc_t\lambda^{1/\rho}}{da_t} = \frac{dE_t[c_{t+1}]}{da_t} \underbrace{E_t[\frac{-u''(c_{t+1})}{-u''(c_t\lambda^{1/\rho})}]}_{\substack{>1 \\ \text{decreasing local convexity}}} + \underbrace{cov_t\left(\frac{dc_{t+1}}{da_t}, \frac{-u''(c_{t+1})}{-u''(c_t\lambda^{1/\rho})}\right)}_{\substack{>0 \\ \text{decreasing consumption risk}}}$$

First, when utility is isoelastic, the local convexity of marginal utility measured by absolute prudence $\frac{u'''(c)}{-u''(c)}$ decreases with consumption c, which is equivalent to saying that the slope of marginal utility $-u''(c)$ is a convex function of marginal utility itself $u'(c)$. As a result, the distributions of current and future consumption that equalize the current and expected marginal utility of consumption do not equalize their slopes: the expected slope is larger than the slope at current consumption $E_t[-u''(c_{t+1})] > -u''(c_t\lambda^{1/\rho})$. It implies that a given increase in future consumption decreases the expected marginal utility of future consumption more than the same increase in current consumption decreases the marginal utility of current consumption: the ratio $E_t[\frac{u''(c_{t+1})}{u''(c_t\lambda^{1/\rho})}]$ is larger than one. Thus, for the current and expected marginal utility to fall by the same amount with a gain in assets, current consumption has to decrease more. Second, when utility is isoelastic and the effect of the shocks on income do not shift direction over time, then consumption is concave in the value of these shocks, as proved in Appendix D: $\frac{d^2c_{t+1}}{da_t\omega^l_{t+1}}$ for any shock l. This means that the response of consumption to a change in assets is decreasing with the value of the transitory and permanent income shocks: consumption responds more in the low consumption states of the world than in the high consumption states. On the contrary, marginal utility is convex so it decrease more with an additional unit of consumption when the transitory and permanent income shocks are small and consumption is low than when the transitory and permanent income shocks are large and consumption is high. The covariance between $\frac{dc_{t+1}}{da_t}$ and $\frac{-u''(c_{t+1})}{-u''(c_t\lambda^{1/\rho})}$ is positive. This raises further the decrease in the expected marginal utility of consumption associated with an additional unit of assets, because an additional unit of assets induces the largest decrease in consumption ($\frac{dc_{t+1}}{da_t}$ large) in the states when a decrease in consumption reduces the marginal utility of future consumption the most ($-u''(c_{t+1})$ large). Because the decrease in expected marginal utility associated with an additional unit of future consumption is reinforced, it means that current consumption must respond even more to a gain in assets than expected future consumption for the current and expected marginal utility to fall by the same amount with such a gain. The precautionary consumption growth between current and expected future consumption shrinks:

$$\frac{d\varphi_t}{da_t} < 0$$

Appendix C Proof of the concavity of consumption in income shocks whose effect does not shift direction

This proof is a generalization of the proof that consumption is concave in assets

Appendix D Departure of the innovation from a random walk

I iterate forward the expression $E_t[c_{t+1}] = c_t\lambda^{1/\rho} + \varphi_t$ and obtain that it implies that $c_t = E_t[c_{t+s}] - \sum_{l=1}^{s} E_t[\varphi_{t+l-1}]$ for all $0 < s < T-t$. I substitute each $E_t[c_{t+s}]$ for $c_t + \sum_{l=1}^{s} E_t[\varphi_{t+l-1}]$ in the intertemporal budget constraint and rearrange to obtain the following equilibrium relationship for consumption:

$$c_{t+1} = \frac{1}{l_{t,T}} \underbrace{\left((1+r)a_t + \sum_{s=0}^{T-t} \frac{E_t[y_{t+s}]}{(1+r)^s}\right)}_{\text{total expected resources}} - \frac{1}{l_{t,T}} \underbrace{\left(\sum_{s=1}^{T-t}\sum_{k=1}^{s} \frac{E_t[\varphi_{t+k-1}]}{(1+r)^s}\right)}_{\substack{\text{total expected} \\ \text{precautionary growth, } > 0}}$$

To implement their precautionary growth, households need to allocate a larger share of their resources to future periods than to the current period. Instead of consuming a fraction $\frac{1}{l_{t,T}}$ of their total expected resources, they first retrieve their expected precautionary consumption growth from their total expected resources and consume a fraction $l_{t,T}$ of what remains. This means that the innovation to consumption depends on both their expected resources and their expected precautionary consumption growth:

$$\xi_{t+1} = \frac{1}{l_{t+1,T}} \sum_{s=0}^{T-t-1} \frac{(E_{t+1}-E_t)[y_{t+1+s}]}{(1+r)^s} - \frac{1}{l_{t,T}} \sum_{s=1}^{T-t-1}\sum_{k=1}^{s} \frac{(E_{t+1}-E_t)[\varphi_{t+k}]}{(1+r)^s}$$

Although the expected value is always zero, the distribution of ξ_{t+1} depends on $(E_{t+1}-E_t)[y_{t+1+s}]$, which in general depends on the level of assets that a household owns. It is therefore not independent from this level of assets though uncorrelated with it.

Appendix E Derivation of the BPP estimating equation

Blundell, Low, and Preston (2013) and Blundell, Pistaferri, and Preston (2008) derive their expression of log-consumption growth by applying three successive approximations, which they assume correspond to the same approximation point: i) at the point where $ln(c_{t+1})$ equals $ln(c_t) + \frac{1}{\rho}ln(\gamma)$ ii) at the point where $ln(c_{t+1})$ equals $E_t[ln(c_{t+1})]$ iii) at the point where the realized income shocks are zero $\varepsilon_{t+1} = 0$ and $\eta_{t+1} = 0$. Applying approximation i) to the Euler equation yields:

$$ln(c_{t+1}) = ln(c_t) + \frac{1}{\rho}ln(\gamma) + \frac{1}{2}E[\tilde{d}(ln(c_{t+1}) - ln(c_t) - \frac{1}{\rho}ln(\gamma))^2] + \mu_{t+1}$$

with $\tilde{d}$ a term that depends on the difference between future and current log-consumption and tends towards a constant when their difference approaches zero, and $\mu_{t+1} = ln(c_{t+1}) - E[ln(c_{t+1})]$ the innovation to log-consumption growth. Assuming that the approximation points i) and ii) coincide, $ln(c_t) + \frac{1}{\rho} ln(\gamma) = E[ln(c_{t+1})]$, and the third term rewrites as $E[n_t \mu_{t+1}^2]$. Applying a Taylor expansion to this term only, it writes as a big O of $E[\mu_{t+1}^2]$. Substituting for it in the expression gives:

$$ln(c_{t+1}) = ln(c_t) + \frac{1}{\rho} ln(\gamma) + O\Big(E[\mu_{t+1}^2]\Big) + \mu_{t+1}$$

Taking an expansion around the point ii) where innovations to log-consumption are zero ($\mu_{t+1} = \mu_{t+2} = ... = \mu_T = 0$), the authors derive an expression of the logarithm of total lifetime consumption:

$$(E_{t+1} - E_t)[ln\left(\sum_{s=0}^{T-t-1} \frac{c_{t+1+s}}{(1+r)^s}\right)] = \mu_{t+1} + O\Big(E[\mu_{t+1}^2], ..., E[\mu_T)^2]\Big)$$

Taking an expansion around the point iii) where income innovations are zero, they derive an expression of the logarithm of total lifetime resources:

$$(E_{t+1} - E_t)[ln\left(\sum_{s=0}^{T-t-1} \frac{y_{t+1+s}}{(1+r)^s} + (1+r)a_{t+1}\right)] = \Phi^\varepsilon \varepsilon_{t+1} + \Phi^\eta \eta_{t+1} + O\Big(E[\varepsilon_{t+1}^2], E[\eta_{t+1}^2], ..., E[\varepsilon_T)^2], E[\eta_T^2]\Big)$$

The intertemporal budget constraint implies that the logarithm of total lifetime consumption coincide with the logarithm of total lifetime resources. Substituting for μ_{t+1} using the previous expression and noting that when the expansion points ii) and iii) coincide, the higher order term $O\Big(E[\mu_{t+1}^2], ..., E[\mu_T)^2]\Big)$ is a $O\Big(E[\varepsilon_{t+1}^2], E[\eta_{t+1}^2], ..., E[\varepsilon_T)^2], E[\eta_T^2]\Big)$, log-consumption growth writes:

$$\Delta ln(c_{t+1}) = \frac{1}{\rho}\beta(1+r) + \Phi^\varepsilon \varepsilon_{t+1} + \Phi^\eta \eta_{t+1} + O\Big(E[\varepsilon_{t+1}^2], E[\eta_{t+1}^2], ..., E[\varepsilon_T)^2], E[\eta_T^2]\Big)$$

In first order approximation around small income shocks, it is:

$$\Delta ln(c_{t+1}) \approx \frac{1}{\rho}\beta(1+r) + \Phi^\varepsilon \varepsilon_{t+1} + \Phi^\eta \eta_{t+1}$$