How Does Consumption Respond to a Transitory Income Shock? Reconciling Natural Experiments and Semi-Structural Estimations

Jeanne Commault*

March 2018

Abstract

Results from natural experiments show that consumption responds strongly and significantly to a transitory variation in income such as a tax rebate or a lottery gain. On the contrary, semi-structural methods that rely on theoretical restrictions to identify in longitudinal survey data the consumption elasticity to a transitory shock find that this elasticity is small and not statistically significant. I show that the two approaches reconcile when relaxing the assumption made by semi-structural methods that the log-consumption growth of a household is independent from the income shocks it has experienced in the past. First, I establish that log-consumption growth depends on past income shocks even in the baseline life-cycle model because of precautionary saving; the correlation is absent from the expression used by semi-structural methods which implicitly assume away precautionary effects. Second, I develop a robust semi-structural estimator that remains consistent in the presence of a correlation between past shocks and log-consumption growth, while previous estimators are downward biased in the presence of such a correlation. Third, I implement the robust estimator in survey data. I find that the elasticity of consumption to a transitory shock is statistically significant and that, on average, 39% of a transitory gain in net income is consumed within the following year, which is comparable to the findings of the literature based on natural experiments.

*Department of Economics, European University Institute, Villa La Fonte, Fiesole, Firenze, Italy, jeanne.commault@gmail.com.

Introduction

How does individual consumption respond to a transitory income shock? The answer has implications for a number of economic questions, including the economy's response to fiscal shocks, the relation between income and consumption inequalities, and the dynamics of business cycles. One obstacle in the way of measuring this answer, however, is that transitory shocks are not usually observed directly. Instead, in longitudinal survey data, households report their total income change, without distinguishing between transitory and permanent changes. Yet, for the researcher, the difference is important because shocks that do not have the same durability should have distinct effects on consumption.

There are two main solutions to overcome this issue, but they yield opposite conclusions. A first approach consists in exploiting specific episodes of observed transitory income variations, such as tax rebates and lottery gains, and pairing them with consumption survey data to directly measure the response of expenditures response to an income shock that the researcher observes and knows to be transitory. In the great majority of these studies, transitory income changes have statistically significant and economically large effects on consumption. [1] A second approach identifies the response of consumption to transitory shocks by putting more structure on the data. Making assumptions about the form of the income process and the way households take their consumption decisions, the seminal paper of Blundell, Pistaferri, and Preston (2008) (hereafter BPP) derives restrictions that can separately identify the elasticity of consumption to transitory and permanent shocks. These assumptions are that income evolves as a transitory-permanent process and that the log-consumption growth of a household does not depend on the income shocks that it has experienced in the past, which is justified as a condition that would hold if the household were to solve a baseline life-cycle model. I refer to the method as a semi-structural estimation because it does not require that the fully-fledged model holds but only that one condition derived from it does. This identification strategy is now influential in all fields of economics that are concerned with the way shocks are passed through decision variables such as consumption and saving but also individual labor supply, family labor supply, time allocation between spouses, and housing or investment. The studies that rely on semi-structural estimation typically find that the elasticity of consumption to a transitory shock is not statistically significant, although it is quite precisely estimated. [2] The divergence between the two approaches means that either the particular transitory shocks observed in natural experiments share specific characteristics that induce households to respond more to them than they do to the shocks they face in survey data, suggesting that these natural experiments have little external validity, or that some of the identifying restrictions imposed by semi-structural methods do not hold well enough in the data to yield reliable estimates.

In this paper, I generalize the semi-structural estimation method by relaxing the assumption that the log-consumption growth of a household does not depend on the income shocks that it has experienced in the past, and find the elasticity of consumption to a transitory shock to be statistically significant and

[1] See for example Souleles (1999), Johnson, Parker, and Souleles (2006), Agarwal, Liu, and Souleles (2007), Parker, Souleles, Johnson, and McClelland (2013), Kaplan and Violante (2014), Misra and Surico (2014) for the consumption response to a tax rebate; Baker and Yannelis (2015), Gelman, Kariv, Shapiro, Silverman, and Tadelis (2015) for the response to the 2013 government shutdown; Agarwal and Qian (2014), Kan, Peng, and Wang (2016) for the response to the distribution of cash or consumption vouchers by the government; Fagereng, Holm, and Natvik (2016) for the response to a lottery gain.

[2] See Appendix A for a review of the literature that builds on the BPP identification strategy. Studies based on longitudinal survey data prior to BPP do not disentangle between transitory and permanent changes, or use specific proxies for permanent and transitory income changes such as disability or short unemployment spells making them close to natural experiments studies (Cochrane (1991), Dynarski and Gruber (1997)).

within the range of values suggested by natural experiments. More precisely, I make three contributions: (i) I establish that, even in a baseline life-cycle model, log-consumption growth depends on the realizations of past shocks, because of precautionary saving that simultaneously depends on a household's history of shocks and affects its log-consumption growth; this precautionary correlation is erroneously assumed away in BPP because of the approximation that the authors use; common extensions to the baseline model generate additional sources of correlation between consumption growth and past shocks; (ii) I develop an estimator of the elasticity of consumption to a transitory shock that is robust to the presence of a correlation between consumption growth and past shocks; existing semi-structural estimators that neglect this correlation are biased downwards if the correlation is negative; (iii) I implement the robust estimator in the same survey data as existing semi-structural estimators and obtain a larger and significant elasticity estimate that is consistent with findings from natural experiments.

First, I prove that, in a baseline life-cycle model, the log-consumption growth of a household is negatively correlated with the transitory shocks that it has experienced in the past. The framework consists of a finite-lived household with isoelastic preferences that maximizes its intertemporal utility subject to a budget constraint and a natural borrowing constraint that never binds—by assumption the household cannot die in debt so it cannot borrow more than the lowest possible amount that it can earn in the future. It faces uncertainty because its future income depends on the realizations of two types of shocks, transitory and permanent.[3] It is the same framework as the one underlying the BPP estimator. The correlation between consumption growth and past income shocks is driven by precautionary behavior, that is, the effect of uncertainty on the decisions of the household. Indeed, when utility is isoelastic, marginal utility is convex, and the presence of uncertainty raises the expected marginal utility of future consumption, inducing the household to shift some of its resources from the present to the future and raising its consumption growth. The magnitude of this precautionary consumption growth depends on the distribution of future consumption, which is influenced by the income shocks that the household has received in the past. In particular, having experienced a good transitory shock in the past raises the amount of assets with which the household enters the current period, shifting upwards the distribution of its future consumption to a region where marginal utility is locally less convex thus where the need for precautionary consumption growth is reduced. The realizations of past transitory income shocks are then negatively correlated with consumption growth.

Semi-structural estimators do not allow for such a correlation: they express the log-consumption growth of a household as a function only of the current income shocks he receives and the changes in its demographic characteristics. The expression is obtained by taking an endogenous approximation of log-consumption growth at the point where the marginal utility of future consumption equals the marginal utility of current consumption, assuming already that future consumption equals current consumption plus a noise of mean zero, thus implicitly assuming away the precautionary trend in consumption.

Generalizing the baseline life-cycle model to incorporate endogenous labor supply, risky assets, uncertainty towards future demographic characteristics, stochastic death, and tighter borrowing constraints does not make the correlation between consumption growth and past income shock disappear. On the contrary, these additional features generate new sources of correlation.

[3] This is consistent with microeconomic data on earnings: Cochrane (1991), Attanasio and Davis (1996), and Brav, Constantinides, and Geczy (2002) show that consumers' revenue is not perfectly insured; the permanent-transitory structure of shocks is found to fit well with the observed dynamics of microeconomic of earnings (MaCurdy (1982), Abowd and Card (1989)).

Second, I develop a robust estimator that allows for of a correlation between log-consumption growth and past shocks. because I relax this assumption, my estimation framework only requires that income evolves as a transitory-permanent income process and remains agnostic about the way households make their consumption decisions. The identification strategy to disentangle the impact of transitory and permanent shocks is to use future income growth as an instrument, because it correlates negatively with the current transitory shock (that reverts to a value of zero) but is independent from the realization of the permanent shock (that stays in the value of log-income for the rest of the household life so it does not affect its income growth after the shock).

When a correlation between log consumption growth and past transitory shocks is present but neglected, it induces researchers to use additional instruments that are in fact endogenous and generate an estimation bias. There exist some periods in the future such that income growth between these periods correlates with the current transitory shock but is independent from past income shocks, but also some periods such that income growth between these periods correlates with both current and past transitory income shocks. Using this income growth as an instrument and failing to disentangle the negative impact of a past transitory shock on current consumption growth from the positive impact of a current transitory shock on current consumption growth leads to underestimating the response of current consumption growth to a contemporaneous transitory shock. The presence of a correlation between log-consumption growth and past shocks makes endogenous also the instruments used to identify the elasticity of consumption to a permanent shock. In that case, however, I am not able to exhibit a robust instrument. The more structural estimators that build on the full life-cycle model to derive an analytical expression of the elasticity of consumption that they can measure neglect precautionary terms as well, so they also yield biased estimates.

Third, I implement the robust estimator in data from the Panel Study of Income Dynamics (PSID) between 1978 and 1992, combined with imputed consumption data from the Consumer Expenditure Survey (CEX) over the same period. This exactly follows BPP, thereby making the comparison between my results and theirs straightforward. With the robust estimator, the elasticity of consumption to transitory shocks is statistically significant. The point estimate of the elasticity shifts to 0.47, which in the data corresponds to an average marginal propensity to consume (MPC) of 0.39%. This says that 39% of a transitory income gain is spent on nondurable consumption within the following year. It is in line with the results from natural experiments, which find that between 9% and 38% of a transitory income gain is used for nondurable consumption expenditures within the following quarter, and between 25% and 60% over the following year although the elasticity is not precisely estimated at the yearly frequency. The elasticity of food consumption to a transitory income shock is statistically significant as well but smaller, and the implied MPC of food consumption is 3%. The elasticity of total consumption, including both nondurable and durable expenditures, is statistically significant and the implied MPC of total consumption is 100%. This is in line with the natural experiment literature as well. When breaking the sample into subgroups with different financial income, wage rate and age, the elasticity to a transitory shock remains large within all subgroups, although larger for households with low financial income, low wage rate, and a young male head, which matches with the results from natural experiments. My findings are robust to making the permanent shocks depreciate over time, allowing households to partly anticipate the shocks, allowing for clustered estimation errors, and having serially correlated measurement error.

Overview

The remainder of the paper is organized as follows. Section 2 exposes the baseline life-cycle model and derives the implications of precautionary behavior in this model. Section 3 describes the identification strategy of the robust estimator, and the bias existing semi-structural estimators presents when log-consumption growth correlates with past shocks. Section 4 details the data and the empirical implementation of the robust estimator. Section 5 presents the results. Section 6 concludes the paper.

1 Correlation between log-consumption growth and past shocks in a life-cycle model

1.1 Household's problem

Time is discrete and indexed by t = 0, 1, ... , T. At each period t, a household i faces the following intertemporal optimization problem:

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

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

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

The household is finite-lived with T the length of its life. It has time-separable preferences. At each period, it derives utility from its consumption expenditures $c_{i,t}$. 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 with a factor β. Preferences are shifted by a vector of demographic variables $z_{i,t}$ whose future values are assumed to be known in advance with certainty by the household. The impact of these demographics on utility is captured by the vector of coefficients δ_t. At period t, the household earns a stochastic amount $y_{i,t}$. The budget constraint (2.3) states that the household has access only to a risk-free asset whose level at the beginning of period t (or equivalently at the end of period $t-1$) I denote $a_{i,t}$ and that delivers the certain interest rate r. The terminal condition on wealth (2.3) states that the household cannot die with a strictly positive level of debt. I assume constancy of the discount factor and of the risk-free interest rate to simplify the presentation but it is straightforward to allow them to be time-varying.

1.2 Income process

The net income of the household i at period t, denoted $y_{i,t}$, is modeled as a transitory-permanent process:

$$ln(y_{i,t}) = p_{i,t} + \mu_{i,t} + f_i + \kappa_t z_{i,t}$$

$$with \quad \begin{cases} p_{i,t} &= p_{i,t-1} + \eta_{i,t} \\ \mu_{i,t} &= \varepsilon_{i,t} + \theta_1 \varepsilon_{i,t-1} + \dots + \theta_q \varepsilon_{i,t-q}. \end{cases}$$

The log of net income is the sum of a permanent component $p_{i,t}$ that follows a random walk, of a transitory component $\mu_{i,t}$ that follows an MA(q) process, and of a term $f_i + \kappa_t z_{i,t}$ capturing an individual fixed effect plus the deterministic influence of demographic variables. The uncertainty of the household comes from the presence of the shocks to the permanent and transitory components of income, $\eta_{i,t}$ and $\varepsilon_{i,t}$. The shock $\eta_{i,t}$ is a permanent shock because its realization remains in the value of p at all following periods, so it affects income until the end of the household's lifetime. On the contrary, the shock $\varepsilon_{i,t}$ is transitory because its effect on income dissipates after q periods, where the order q is to be established empirically (at the end of the household's life, after period $T-q$, the effect of a transitory shock lasts until period T, that is, until the end of the household's lifetime T so the transitory shock resembles a permanent shock). The permanent and transitory shocks are drawn independently from each other and from their previous realizations at each period. They are normalized to have mean zero. The distributions from which they are drawn may vary over time and across subgroups of households.

The demographic variables $z_{i,t}$ are assumed not to be subject to any uncertainty: they may change over time, but these variations are expected by consumers. Their impact on utility is measured by the vector κ_t, which is allowed to shift with calendar time.

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

1.3 Precautionary behavior

My demonstration that there exists a strictly negative correlation between past transitory shocks and log-consumption growth is in two steps: (i) precautionary behavior raises expected consumption growth; (ii) a positive past transitory shock, equivalent to gain in current net assets, reduces the need for precautionary behavior, thus reduces expected consumption growth.

I start with point (i). The equilibrium condition of the households' problem, known as the Euler equation, is as follows:

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

It states that the household aims at equalizing its expected marginal utility over time. 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.[4] I introduce the equivalent precautionary premium associated with consumption uncertainty at $t+1$, which is the consumption loss φ_t that a household facing no uncertainty would have to incur for its marginal utility to be as high as it is in the presence of uncertainty:

$$E_t[u'(c_{t+1}R_{t,t+1}^{-1/\rho})] = u'(E_t[c_{t+1}]R_{t,t+1}^{-1/\rho} - \varphi_t).$$

Kimball (1990b) initially develops the notion of equivalent precautionary premium to compare consumption and wealth in a certain and in an uncertain environment at a given level of expected marginal utility. I show that it also measures the difference between a household's current (certain) consumption and its future (uncertain) consumption, which have to yield the same expected marginal utility. Substituting for

[4]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.

$E_t[u'(c_{t+1}R_{t,t+1}^{-1/\rho})]$ in the Euler equation, I obtain:

$$
\begin{aligned}
u'(c_t) &= u'(E_t[c_{t+1}]R_{t,t+1}^{-1/\rho} - \varphi_t) \\
c_t &= (E_t[c_{t+1}]R_{t,t+1}^{-1/\rho} - \varphi_t) \\
E_t[c_{t+1}] &= c_t + \underbrace{(R_{t,t+1}^{1/\rho} - 1)c_t}_{\text{int. subs. + dem.}} + \underbrace{\varphi_t R_{t,t+1}^{1/\rho}}_{\text{precaution}}.
\end{aligned}
\tag{2.4}
$$

Under the joint assumption of a risky income path and a strictly convex marginal utility (a property verified by isoelastic preferences), Jensen's inequality implies that φ_t is strictly positive.[5] Thus, the presence of uncertainty induces the household to set its current consumption below its future expected consumption by an amount φ_t. Intuitively, when the household's marginal utility is convex, the impact of a loss and of a gain in consumption are asymmetric: marginal utility increases more following a loss than it decreases following a gain. Thus, on average, the possibility of consumption losses dominates and the presence of uncertainty towards future consumption, that is the presence of mean-zero shocks to future consumption, raises the expected marginal utility of future consumption. As a result, in the presence of uncertainty, an optimizing household moves additional resources from the present to the future, which is associated with a higher expected marginal utility. Under perfect foresight or under certainty equivalence $\varphi_t = 0$ because either there are no shocks or their impact on future marginal utility is symmetric and averages out. Intertemporal substitution and changes in demographics also affect expected consumption growth, modifying it by an amount $(R_{t,t+1}^{1/\rho} - 1)c_t$, so their effect is proportional to current consumption.

Figures 1.(a) and 1.(b) present a graphical illustration of these mechanisms. The top figure 1.(a) pictures a situation in which future consumption can take two values, a low value denoted c_{t+1}^L, and a high value, denoted c_{t+1}^H. 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, as pictured on the bottom figure 1.(b): instead of having a flat expected path (with actual consumption changes being driven exclusively by income news), a household that faces uncertainty allocates more consumption to the future periods.

To derive an expression for log-consumption growth, I take the relation (2.4), substitute for $E_t[c_{t+1}] = c_{t+1} - \xi_{t+1}$ with ξ_{t+1} the innovation to future consumption, divide each side by $c_t R_{t,t+1}^{1/\rho}$ and take the logarithm. I then apply a Taylor expansion around the point where income shocks are equal to zero: $(\varepsilon_{t+1}, \eta_{t+1}) = (0,0)$, denoting with an exponent 0 the value of the variables taken at this point. It yields

[5] When marginal utility $u'(c)$ is strictly convex, Jensen's inequality states that:

$$E_t\left[u'(c_{t+1})\right] > u'(E_t[c_{t+1}]) \;\Leftrightarrow\; u'(E_t[c_{t+1}] - \varphi_t) > u'(E_t[c_{t+1}]) \;\Leftrightarrow\; E_t[c_{t+1}] - \varphi_t < E_t[c_{t+1}] \;\Leftrightarrow\; 0 < \varphi_t.$$

Figure 1

(a) Equalization of expected marginal utility

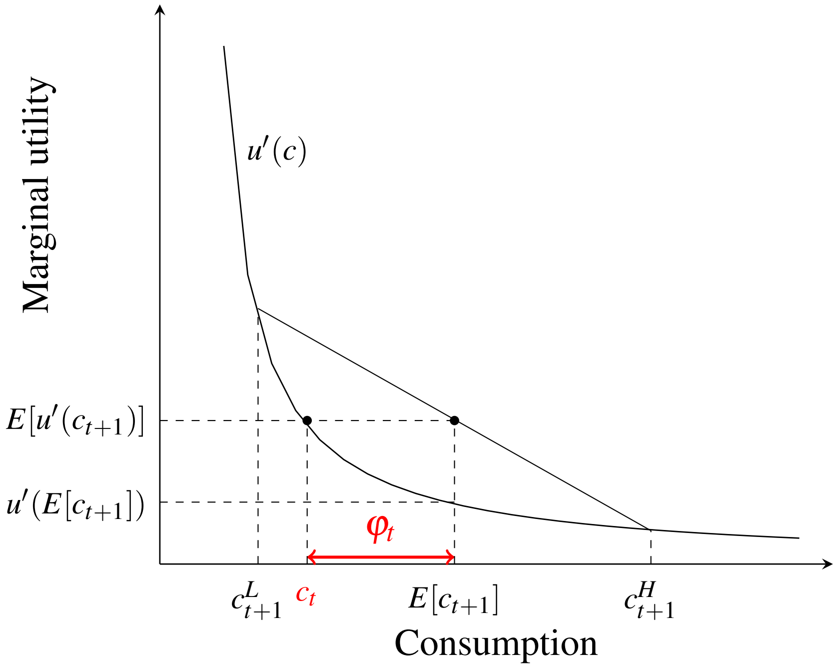

(b) Expected consumption path

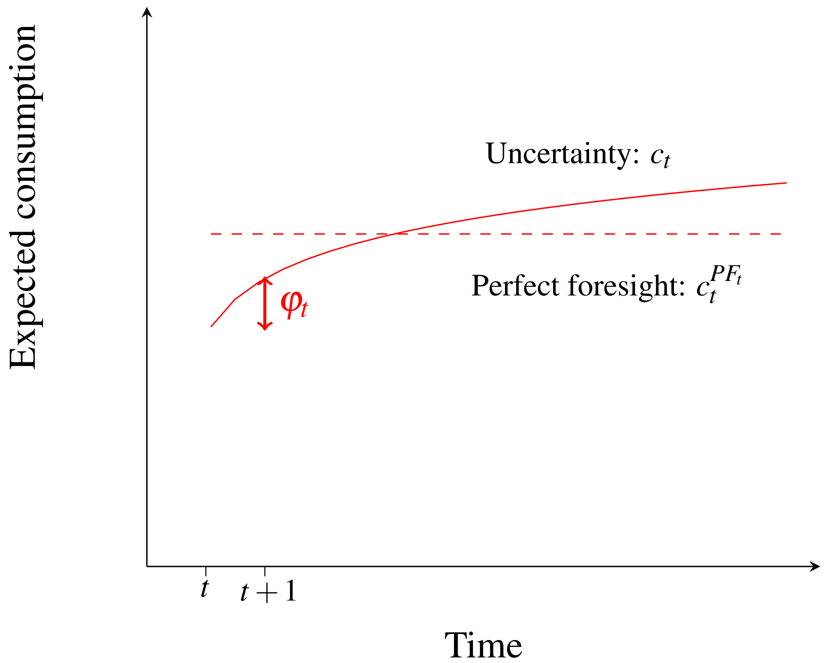

Figure 1.(a) pictures how current and future consumption are set, so that the marginal utility of current consumption be equal to the expected marginal of future consumption. Figure 1.1 (b) represents the expected consumption path of the household. These graphs are only illustrative.

the following decomposition of log-consumption growth:

$$\Delta ln(c_{t+1}) = \underbrace{\frac{1}{\rho} ln(\beta(1+r))}_{\text{int. substitution}} + \underbrace{\frac{1}{\rho}\Delta(\delta_{t+1} z_{t+1})}_{\text{demographics}} + \underbrace{ln(1+\frac{\varphi_t + \xi^0_{t+1}}{R^{1/\rho}_{t,t+1} c_t})}_{\text{precaution}}$$
$$+ \varepsilon_{t+1} \underbrace{\left(\frac{dln(c_{t+1})}{d\varepsilon_{t+1}}\right)^0}_{\substack{\text{elasticity to} \\ \text{transitory shocks: } \phi^\varepsilon_t}} + \eta_{t+1} \underbrace{\left(\frac{dln(c_{t+1})}{d\eta_{t+1}}\right)^0}_{\substack{\text{elasticity to} \\ \text{permanent shocks: } \phi^\eta_t}} + o(\varepsilon_{t+1}, \eta_{t+1}) \qquad (2.5)$$

The first three terms correspond to an approximation of expected log-consumption growth, and the next two form an approximation of the innovation to log-consumption growth. Because the presence of uncertainty raises expected consumption growth by φ_t, it increases expected log-consumption growth by $ln(1+\frac{\varphi_t+\xi^0_{t+1}}{R^{1/\rho}_{t,t+1}c_t})$. The innovation to future consumption ξ^0_{t+1} captures precautionary behavior too because it would be zero under perfect foresight. The only reason it is not zero in general is that, because of future precautionary consumption growth, consumption c_{t+1} is not linear in income but strictly concave, so that consumption innovation is strictly positive at the point where income innovations are equal to their expected value: $c^0_{t+1} = c_{t+1}(0,0) = c_{t+1}(E_t[\varepsilon_{t+1}], E_t[\eta_{t+1}]) > E_t[c_{t+1}]$.

Let me now turn to point (ii) and show that this precautionary component of expected log-consumption growth is negatively correlated with past shocks.

Theorem: In the model presented above, the equivalent precautionary premium, φ_t, is negatively correlated with net assets and with the value of past transitory income shocks. At any period $0 < t < T$, and for any $0 < k < t$:

$$\frac{d\varphi_t}{da_t} < 0 \text{ and } \frac{d\varphi_t}{d\varepsilon_{t-k}} < 0.$$

Intuition: 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:

$$\varphi_t = \frac{1}{2} \underbrace{p(E_t[c_{t+1}])}_{\text{absolute prudence}} \times \underbrace{var_t(c_{t+1})}_{\text{consumption risk}} + o(var_t(c_{t+1})).$$

The coefficient of absolute prudence measures the local curvature of the marginal utility function at the point $c = E_t[c_{t+1}]$. 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 variance of consumption summarizes the magnitude of consumption fluctuations.

I take the derivative with respect to net assets of this approximation of φ_t:

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

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 condenses its distribution. Indeed, a gain in assets does not affect consumption identically in all states of the world but raises more consumption in the low consumption states of the world and raises it less 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. The exact proof is presented in Appendix B.

Figures 2.(a) and 2.(b) illustrate this graphically. The top figure 2.(a) represents the effect of a shift upwards in future consumption around a higher expected value. Because the convexity is less pronounced around higher levels of consumption, the precautionary growth necessary to equalize the marginal utility of current and future consumption decreases.[6] The bottom figure 2.(b) pictures the consequence of an assets gain for the expected consumption path. Under perfect foresight (dashed line), a gain in assets raises current and expected consumption by exactly the same amount: the expected consumption path under perfect foresight is simply translated upwards from it initial position (in red) to its final position (in blue) but the slope remains unchanged. In the presence of uncertainty (plain line), the increase in assets also modifies the slope of the expected consumption path, and not only its level, because the precautionary component of the slope φ depends on the assets that a household owns. A gain in assets reduces the need for precautionary growth, so the expected consumption path becomes less upward-sloping, in addition to being shifted upwards.

Now, changes in the value of the past transitory shocks experienced by the household are equivalent to changes in the amount of net assets wth which it enters the period. This is because the realizations of past transitory shocks (prior to the immediate past period) have no impact on the determinants of current consumption other than on net assets. For any $k > 1$:

$$\frac{d\varphi_t}{d\varepsilon_{t-k}} = \underbrace{\frac{da_t}{d\varepsilon_{t-k}}}_{>0}\underbrace{\frac{d\varphi_t}{da_t}}_{<0} < 0.$$

[6]For the sake of simplicity, I plotted the consequence of an homogeneous shift in consumption in this figure, by which future consumption increases by the same amount in the low state and high state of the world. Because consumption is concave in assets, a gain in assets increases in fact its consumption more in the low state than in the high state of the world. This means that not only does the distribution of future consumption shift up, but the gap between consumption in the low and high states of the world decrease, reducing further the need for precautionary consumption growth.

Figure 2

(a) Precautionary consumption growth before (in red) and after (in blue) a gain in assets

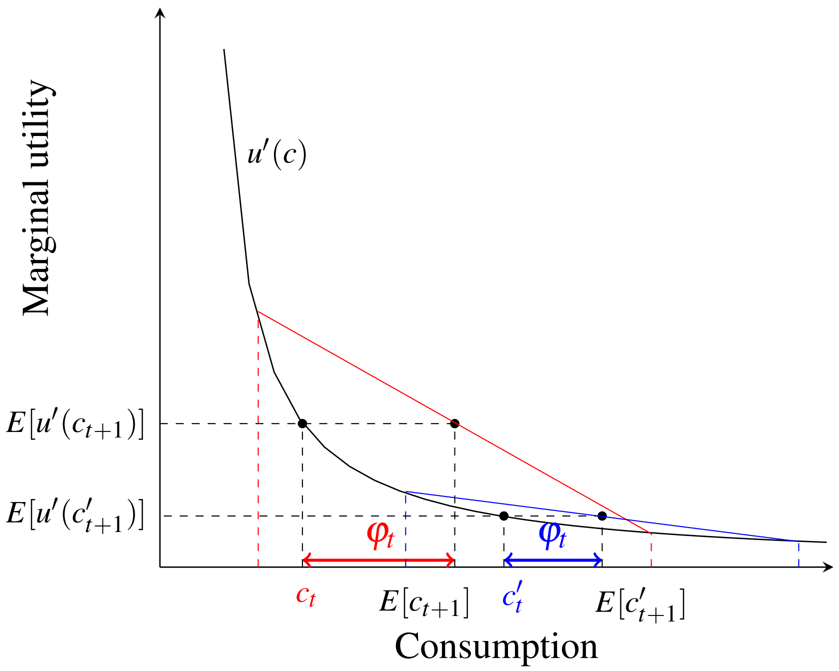

(b) Expected consumption path before (in red) and after (in blue) a gain in assets

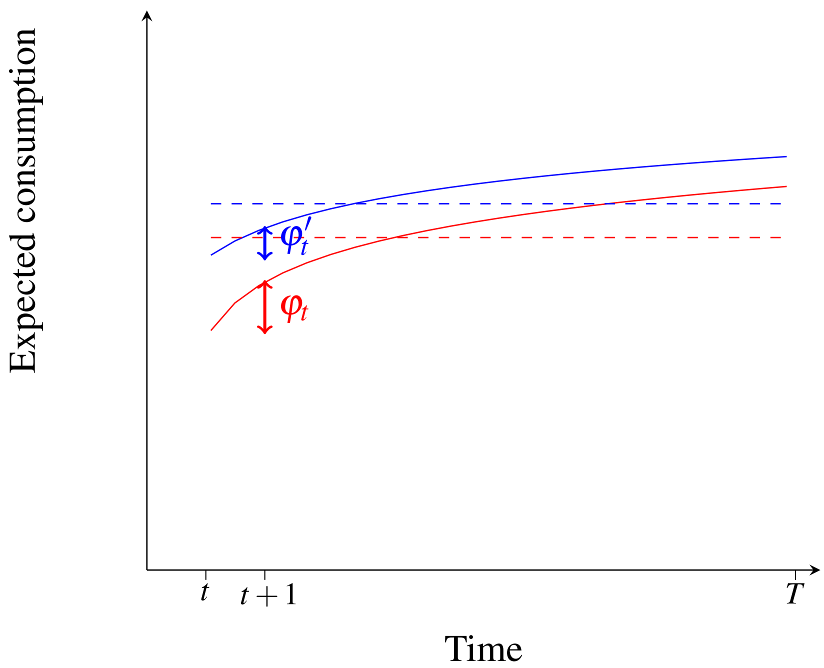

Figure 1.2 (a) pictures the change in the choice of current consumption, following a determinate increase in future consumption (a shift upwards). Figure 1.2 (b) represents a related shift in the expected consumption path of the household, following a same gain in net assets. These graphs are only illustrative.

Indeed, since they are transitory they do not affect the distribution of future income after $q = 1$ periods. Their impact on net assets is positive because a transitory gain is smoothed over the household's lifetime, thus passed on to later periods through increased assets: $\frac{da_t}{d\varepsilon_{t-k}} > 0$.

A transitory shock at $t-1$ impacts both a_t and y_t but not the distribution of income at $t+1$ and later, so the same intuition applies[7]

$$\frac{d\varphi_t}{d\varepsilon_{t-1}} = \underbrace{\frac{da_t}{d\varepsilon_{t-1}}}_{>0} \underbrace{\frac{d\varphi_t}{da_t}}_{<0} + \underbrace{\frac{dy_t}{d\varepsilon_{t-1}}}_{>0} \underbrace{\frac{d\varphi_t}{dy_t}}_{<0} < 0.$$

A detailed demonstration is provided in Appendix B.

The Theorem implies that log-consumption growth is strictly decreasing in the value of past transitory shocks. Indeed, it proves that $\frac{d\varphi_t}{d\varepsilon_{t-k}} < 0$. In addition, I show that $\frac{d\xi^0_{t+1}}{d\varepsilon_{t-k}} < 0$, because $\frac{dc_{t+1}}{da_t}$ is convex in ε_{t+1} and η_{t+1}.[8] On the contrary, consumption c_t is positively correlated with ε_{t-k}. The terms ρ, β, r, δ_{t+1} and z_{t+1} are exogenous terms that do not respond to income shocks. As a result $cov_t(ln(1+\frac{\varphi_t+\xi^*_{t+1}}{R^{1/\rho}_{t,t+1}c_t}), \varepsilon_{t-k}) < 0$ and:

$$cov_t(\Delta ln(c_{t+1}), \varepsilon_{t-k}) = \underbrace{cov_t(\frac{1}{\rho}ln(\beta(1+r)) + \frac{1}{\rho}\delta_{t+1}z_{t+1}, \varepsilon_{t-k})}_{=0} + \underbrace{cov_t(ln(1+\frac{\varphi_t+\xi^0_{t+1}}{R^{1/\rho}_{t,t+1}c_t}), \varepsilon_{t-k})}_{<0}$$
$$+ \underbrace{E_t[\varepsilon_{t+1}]}_{=0} cov_t(\phi^\varepsilon_t, \varepsilon_{t-k}) + \underbrace{E_t[\eta_{t+1}]}_{=0} cov_t(\phi^\eta_t, \varepsilon_{t-k})$$
$$cov_t(\Delta ln(c_{t+1}), \varepsilon_{t-k}) < 0.$$

1.4 Approximation of consumption growth in semi-structural estimators

Semi-structural estimators approximate the log-consumption growth of a household i as a random walk:

$$\Delta ln(c^{BPP}_{i,t+1}) \approx \underbrace{\Gamma_{t+1}}_{\text{int. sub. + dem. + prec.}} + \varepsilon_{i,t+1}\phi^\varepsilon_{i,t} + \eta_{i,t+1}\phi^\eta_{i,t} \tag{2.6}$$

with Γ_{t+1} a deterministic trend, constant across households with the same demographics. Expected log-consumption growth, which coincides with Γ_{t+1}, is therefore independent from each household's level of assets: two households with the same demographics but different levels of assets have the same trend. This implies that the precautionary component of expected log-consumption growth is deterministic (because the trend is and all the other components of the trend are exogenous parameters). Therefore, the expression allows the household to engage in precautionary behavior but only in a deterministic way. The household knows at the beginning of his life how much it will save for precautionary reasons at each period and cannot revise its precautionary saving upon observing the income shocks it receives. Also, precautionary saving must be the same for all households who have the same demographic characteristics. Thus, precautionary behavior has the same impact as the demographic characteristics. The expression imposes that the covariance between log-consumption growth and past shocks is zero. For

[7] The notation $\frac{d\varphi_t}{dy_t}$ is a little imprecise: it refers to a change in income y_t that has no impact on income at later periods.

[8] It implies that $\frac{dc_{t+1}}{da_t}(E_t[\varepsilon_{t+1}], E_t[\eta_{t+1}]) < \frac{dE_t[c_{t+1}]}{da_t}(\varepsilon_{t+1}, \eta_{t+1})$.

any $k > 0$:

$$cov_t(\Delta ln(c^{BPP}_{i,t+1}), \varepsilon_{i,t-k}) = \underbrace{cov_t(\Gamma_{i,t+1}, \varepsilon_{i,t-k})}_{=0} + \underbrace{E_t[\varepsilon_{i,t+1}]}_{=0} cov_t(\phi^{\varepsilon}_{i,t}, \varepsilon_{i,t-k}) + \underbrace{E_t[\eta_{i,t+1}]}_{=0} cov_t(\phi^{\eta}_{i,t}, \varepsilon_{i,t-k}) = 0.$$

Assuming away precautionary behavior How come that the BPP approximation around small income innovations derived from this same model obtains an expression (i) in which log-consumption growth is independent from the realizations of past shocks, and (ii) in which the consumption elasticity to income shocks is the percentage change in total expected resources, as under perfect foresight? I show this is because, in their Taylor approximation of the Euler equation, BPP assimilate the gap between actual log-consumption and the value of log-consumption that would be chosen under perfect foresight:

$$e_{t+1} = \ln(c_{t+1}) - (ln(c_t) + ln(R^{1/\rho}_{t,t+1})),$$

and the innovation to log-consumption growth:

$$u_{t+1} = \ln(c_{t+1}) - E_t[ln(c_{t+1})]$$

$$u_{t+1} = \ln(c_{t+1}) - (ln(c_t) + ln(R^{1/\rho}_{t,t+1}) + E_t[ln(1 + \frac{\varphi_t + \xi_{t+1}}{R^{1/\rho}_{t,t+1}})])$$

$$u_{t+1} = e_{t+1} - E_t[ln(1 + \frac{\varphi_t + \xi_{t+1}}{R^{1/\rho}_{t,t+1}})].$$

They take at the same time an approximation around $e_{t+1} = 0$ and $u_{t+1} = 0$. As a result, they are neglecting the contribution of precautionary behavior, $E_t[ln(1 + \frac{\varphi_t + \xi_{t+1}}{R^{1/\rho}_{t,t+1}})]$. I detail their computations in Appendix B.

1.5 Generalization

I consider a more comprehensive model in which the household i faces the following intertemporal optimization problem:

$$\max_{c_{i,t}, c^D_{i,t}, h_{i,t}, \ldots} \sum_{s=0}^{T-t} E_t \left[\left(u(c_{i,t+s}) + u^D(c^D_{i,t+s}) + v(h_{i,t+s}) \right) \left(\Pi^s_{l=0} \beta^{t+l} \right) e^{\delta_{t+s} z_{i,t+s}} | \Pi^s_{l=0} L_{t+l} = 1 \right] \quad (2.3)$$

$$s.t. \quad \left(a_{k+1} + a^r_{k+1} \right) \mathbb{1}_{L_{t+1}=1} = (1+r)a_k + (1+r^r_k)a^r_k - c_k + h_k w_k \qquad \forall t \leq k \leq T, \quad (2.4)$$

$$\left(a_t + a^r_t \right) \mathbb{1}_{p_{t+1} \neq 1} \geq 0 \quad (2.3)$$

The household is finite-lived with T the maximum length of its life. It has time-separable preferences. At each period, it derives separately utility from its non-durable consumption expenditures $c_{i,t}$, utility its durable consumption expenditures $c^D_{i,t}$, and disutility from the hours of work of its household members $h_{i,t}$. The period utility functions $u^{ND}(c)$ is isoelastic and its functional form is $u(c) = \frac{c^{1-\rho}-1}{1-\rho}$. Future utility and disutility are discounted at period t with a factor β_t. Preferences are shifted by a vector of demographic variables $z_{i,t}$ whose values can be stochastic. The impact of these demographics on utility is captured by the vector of coefficients δ_t. The household has a probability p_t to be alive at the beginning of period t, in which case the life dummy variable $L_t = 1$. Each hour of work at period t is paid at a stochastic wage-rate $w_{i,t}$. The budget constraint (2.3) states that the household has access to a risk-free asset whose amount at the beginning of period t I denote $a_{i,t}$ and that delivers the certain interest rate r, and a risky asset whose amount at the beginning of period t I denote $a^r_{i,t}$ and that delivers an uncertain interest rate r^r_t. The terminal condition on wealth (2.3) states that the household cannot die with strictly positive amounts of debt: when the probability of being alive at the next period is not one, the household cannot have strictly positive amounts of debt.

The wage-rate follows the same transitory-permanent process as earnings :

$$
\begin{aligned}
&ln(w_{i,t}) = p_{i,t} + \mu_{i,t} + fe_i + \kappa_t z_{i,t} \\
&with \begin{cases} p_{i,t} &= p_{i,t-1} + \eta_{i,t} \\ \mu_{i,t} &= \varepsilon_{i,t} + \theta_1 \varepsilon_{i,t-1} + ... + \theta_q \varepsilon_{i,t-q}. \end{cases}
\end{aligned}
$$

with $p_{i,t}$ the permanent component of the wage rate, $\mu_{i,t}$ its transitory component, fe_i a fixed-effect, and $\kappa_t z_{i,t}$ a term measuring the impact of demographic characteristics. The permanent component evolves as a random walk with innovation $\eta_{i,t}$ called the permanent shock, and the permanent component as an MA(q) with innovation $\varepsilon_{i,t}$ called the transitory shock. The realizations of the shocks are non-serially correlated, non-correlated between the permanent and the transitory shock, and orthogonal to any past variable.

One first order condition associated with this maximization problem, known as the risk-free Euler equation is:

$$
\begin{aligned}
&u'(c_t) = E_t[u'(c_{t+1})\beta_{t+1}(1+r)e^{\Delta\delta_{t+1}z_{t+1}}]p_{t+1} \\
&u'(c_t) = E_t[u'(c_{t+1}\tilde{R}^{-1/\rho}_{t,t+1})], \quad \text{with} \quad \tilde{R}_{t,t+1} = \beta_{t+1}(1+r)e^{\Delta\delta_{t+1}z_{t+1}}p_{t+1}
\end{aligned}
$$

The household seeks to equalize its marginal utility of consumption over time, with its future marginal utility weighted by its future discount factor, the risk-free interest factor, the change in its demographic characteristics, and its probability of still being alive in the future. Because utility is isoelastic, these effect can be integrated in the argument of the utility function as a weight $\tilde{R}^{-1/\rho}_{t,t+1}$ on future consumption. Marginal utility is convex, so the expected value of the marginal utility of weighted future consumption is larger than the marginal utility of expected weighted future consumption, and current consumption is smaller than expected weighted future consumption. I denote φ_t the difference between current nondurable consumption and expected weighted future nondurable consumption. It now incorporates the effect of the uncertainty about c_{t+1} but also that of the uncertainty about $\tilde{R}^{-1/\rho}_{t,t+1}$. Expected consumption growth writes:

$$
E_t[c_{t+1}] = c_t + \underbrace{(E_t[\tilde{R}^{-1/\rho}_{t,t+1}]^{-1} - 1)c_t}_{\text{int. subs. + dem.}} + \underbrace{cov(c_{t+1}, \tilde{R}^{-1/\rho}_{t,t+1})}_{\text{interactions with cons.}} + \underbrace{\varphi_t E_t[\tilde{R}^{-1/\rho}_{t,t+1}]^{-1}}_{\text{precaution}}. \tag{2.4}
$$

Consumption growth incorporates an additional term that captures the fact that the expected product of future consumption and the weight $\tilde{R}_{t,t+1}^{-1/\rho}$ is not the product of their expected values because they can covary. The precautionary component does not disappear when generalizing the model. The term $\varphi_t E_t[\tilde{R}_{t,t+1}^{-1/\rho}]^{-1}$ depends on the convexity of $u'(.)$ and on the distribution of $c_{t+1}\tilde{R}_{t,t+1}^{-1/\rho}$ that is expected at period t. The value of $c_{t+1}\tilde{R}_{t,t+1}^{-1/\rho}$ is uncertain at t because of the household faces shocks to its future income, but also because of its probability of dying, the riskiness of its portfolio, and the shocks to its future demographics and discount factor. Therefore, there would be a precautionary component of consumption growth even in the absence of uncertainty about future income. As in the baseline model, current variables such as the level of risk-free assets that the household owns influence the distribution of its future consumption. The realizations of past transitory shocks on its wage-rate determine the amount of assets owned by a household, so they correlate with its consumption growth through this precautionary component. I do not prove, however, that the level of assets is strictly negatively correlated with precautionary saving: in this more general model, it could be for example that experiencing a good transitory shock in the past induces the household to invest more heavily in risky assets, raising its need for precautionary consumption growth. As a result, except in the knife-edge case where the effects of a past transitory shock on φ_t through risk-free assets, risky assets, hours worked, and current income compensate exactly, log-consumption growth varies with the value of the past transitory income shock:

$$\frac{d\Delta ln(c_{t+1})}{d\varepsilon_{t-s}} \neq 0$$

2 Model and robust estimator

2.1 Statistical model

Because the life-cycle model is the standard framework in the consumption literature, it is important that the estimating restrictions used are consistent with this framework, but the estimation of consumption elasticities does not require imposing that the life-cycle model holds. The estimation framework I consider allows for more flexible behaviors than the ones dictated by a life-cycle model, more flexible even than the behaviors allowed by a generalized life-cycle model. I refer to this set of restrictions as the statistical model.

Log-consumption growth The log-consumption growth of a household is modeled as a flexible function of the current and past income shocks it has received. I take an exact Taylor expansion around the point where the realizations of the current transitory and permanent income shocks are equal to zero:

$$\Delta(ln(c_{i,t}) - \delta_t z_{i,t}) = f_{i,t}(\varepsilon_{i,t}, \eta_{i,t}, \varepsilon_{i,t-1}, \eta_{i,t-1}, ., \zeta_{i,t}, \zeta_{i,t-1})$$

$$\Delta(ln(c_{i,t}) - \delta_t z_{i,t}) = \underbrace{f_{i,t}(0, 0, \varepsilon_{i,t-1}, \eta_{i,t-1}, ., \zeta_{i,t}, \zeta_{i,t-1})}_{\text{uncorrelated with current income shocks}} + \sum_{s=1}^{\infty} \frac{\varepsilon_{i,t}^s}{s!}\left(\frac{d^s\Delta ln(c_{i,t})}{d\varepsilon_{i,t}^s}\right)^0 + \frac{\eta_{i,t}^s}{s!}\left(\frac{d^s\Delta ln(c_{i,t})}{d\eta_{i,t}^s}\right)^0$$

Log-consumption growth is a function of the current and past income shocks received by the household that can be non-linear, household and period specific, and can depend on other shocks ζ. These shocks include measurement error for instance. The income shocks are assumed to be non-serially correlated, the other shocks to be uncorrelated with the income shocks, and the function $f_{i,t}(.)$ to be independent from the realizations shocks whose effect it models, so the zeroth-order term of the expansion is uncorrelated with the current income shocks. The first order terms measure the effect of the transitory and permanent shocks on log-consumption growth, thus their effect on log-consumption since shocks are uncorrelated with past variables. The coefficients therefore coincide with the elasticity of consumption to the transitory and permanent shocks around the point where the shocks are small. The remainder from the Taylor expansion is $o(\varepsilon_{i,t}, \eta_{i,t}) = \sum_{s=2}^{\infty} \frac{\varepsilon_{i,t}}{s!} \left(\frac{d^s \Delta ln(c_{i,t})}{d\varepsilon_{i,t}^s}\right)^0 + \sum_{s=2}^{\infty} \frac{\eta_{i,t}}{s!} \left(\frac{d^s \Delta ln(c_{i,t})}{d\eta_{i,t}^s}\right)^0$. Demographic characteristics z are allowed to be influenced by the realizations of the current income shocks. I remove the effect of these demographic characteristics and consider detrended consumption. This shifts the interpretation of the elasticity from $\frac{d\Delta ln(c_{i,t})}{d\varepsilon_{i,t}}$ to $\frac{d\Delta(ln(c_{i,t})-\delta_t z_{i,t})}{d\varepsilon_{i,t}}$: it removes from the response of log-consumption to an income shock the part of the response that is driven by the indirect effect of the shock through changes in demographic characteristics. For instance, finding a job after a short period of unemployment increases the revenue of the household but being employed commands some specific consumption expenditures: detrending log-consumption growth from the employment status of the head removes from the response of consumption these expenditures associated with being employed. To be able to do that, however, I need to assume that current log-consumption growth is linearly related with the current change in demographics. An advantage of detrending is to make the estimating moments robust to the presence of a serial correlation in the changes in demographics experienced by a household. Another advantage is to make the elasticities more similar over time and across households, and therefore to gain precision in the estimation.

Log-income growth The log-income of the household is a transitory-permanent income process, which implies the following expression for its growth:

$$\Delta ln(y_{i,t}) - \kappa_t z_{i,t} = \eta_{i,t} + \varepsilon_{i,t} - (1-\theta)\varepsilon_{i,t-1} - \theta\varepsilon_{i,t-2} + \mu_{i,t}^y - \mu_{i,t-1}^y$$

Log-income growth is a linear function of the current permanent shock, and the current and past transitory shocks up to period $t-q-1$, where q is order of the $MA(q)$ process that the transitory component follows. This order is determined empirically. In the dataset I consider, because the covariance of log-income growth with future log-income growth is no longer statistically significant after $t+2$, I set $q=1$ and denote $\theta_1 = \theta$. Because demographic characteristics z affect current log-income, the current change in demographic characteristics enters the expression of current log-income growth. Because measurement error μ^y can affect log-income growth, the change in measurement error the expression of current log-income growth as well. Measurement error is assumed to be classical, which means that it is serially uncorrelated and uncorrelated with potential measurement error in consumption.

Relations between households All shocks are drawn independently across households, except for shocks affecting household-period observations that belong to the same cluster, where a cluster is defined empirically.

Average elasticity The individual elasticities of consumption to income shocks are:

$$\phi^{\varepsilon}_{i,t} = \frac{d\Delta(ln(c_{i,t}) - \delta_t z_{i,t})}{d\varepsilon_{i,t}} = \sum_{s=1}^{\infty} \frac{\varepsilon^{s-1}_{i,t}}{(s-1)!} \left(\frac{d^s \Delta ln(c_{i,t})}{d\varepsilon^s_{i,t}} \right)^0$$

$$\phi^{\eta}_{i,t} = \frac{d\Delta(ln(c_{i,t}) - \delta_t z_{i,t})}{d\eta i,t} = \sum_{s=1}^{\infty} \frac{\eta^{s-1}_{i,t}}{(s-1)!} \left(\frac{d^s \Delta ln(c_{i,t})}{d\eta^s_{i,t}} \right)^0$$

They writes as weighted sums of the derivatives of log-consumption around small income shocks. I denote $\phi^{\varepsilon} = E[\phi^{\varepsilon}_{i,t}]$ and $\phi^{\eta} = E[\phi^{\eta}_{i,t}]$ the average values of the consumption elasticities to transitory and permanent income shocks in the population over the sample period (without household and time indices since they represent average values). Their expressions are:

$$E[\phi^{\varepsilon}_{i,t}] = \sum_{s=1}^{\infty} \frac{E[\varepsilon^{s-1}_{i,t}]}{(s-1)!} E\left[\left(\frac{d^s \Delta ln(c_{i,t})}{d\varepsilon^s_{i,t}} \right)^0 \right] = \sum_{k=1}^{\infty} \frac{(2k-3)!!E[\varepsilon^2_{i,t}]^{k-1}}{(2k-2)!} E\left[\left(\frac{d^{2k-1} \Delta ln(c_{i,t})}{d\varepsilon^{2k-1}_{i,t}} \right)^0 \right]$$

$$E[\phi^{\eta}_{i,t}] = \sum_{s=1}^{\infty} \frac{E[\eta^{s-1}_{i,t}]}{(s-1)!} E\left[\left(\frac{d^s \Delta ln(c_{i,t})}{d\eta^s_{i,t}} \right)^0 \right] = \sum_{k=1}^{\infty} \frac{(2k-3)!!E[\eta^2_{i,t}]^{k-1}}{(2k-2)!} E\left[\left(\frac{d^{2k-1} \Delta ln(c_{i,t})}{d\eta^{2k-1}_{i,t}} \right)^0 \right]$$

By construction, the realizations of the current income shocks are uncorrelated with the derivative of log-consumption at the point where the current income shocks are zero. Also, because shocks are drawn from normal distributions, the moments of odd orders are zero so that only the moments of even order remain and they are all multiples of the second order moment $E[\varepsilon^{2k}_{i,t}] = (2k+1)!!E[\varepsilon^2_{i,t}]^k$. Now, because shocks are normal, the covariance of log-consumption growth with an income shock is a multiple of the second order moment of the distribution from which the shock is drawn. A consistent estimator of the average elasticity to an income shock in the population is then the ratio of the covariance between log-income growth and the shock over the variance of the shock:

$$\begin{aligned}
\hat{\phi}^{\varepsilon} = \frac{cov(\Delta ln(c_{i,t}), \varepsilon_{i,t})}{var(\varepsilon_{i,t})} &= \frac{1}{E[\varepsilon^2_{i,t}]} E\left[\sum_{s=1}^{\infty} \frac{\varepsilon^{s+1}_{i,t}}{s!} \left(\frac{d^s \Delta ln(c_{i,t})}{d\varepsilon^s_{i,t}} \right)^0 \right] \\
&= \frac{1}{E[\varepsilon^2_{i,t}]} \sum_{k=1}^{\infty} \frac{(2k-1)!!E[\varepsilon^2_{i,t}]^k}{(2k-1)!} E\left[\left(\frac{d^{2k-1} \Delta ln(c_{i,t})}{d\varepsilon^{2k-1}_{i,t}} \right)^0 \right] \\
&= \sum_{k=1}^{\infty} \frac{(2k-3)!!E[\varepsilon^2_{i,t}]^{k-1}}{(2k-2)!} E\left[\left(\frac{d^{2k-1} \Delta ln(c_{i,t})}{d\varepsilon^{2k-1}_{i,t}} \right)^0 \right] = \underbrace{E[\phi^{\varepsilon}_{i,t}]}_{\phi^{\varepsilon}} \\
\hat{\phi}^{\eta} = \frac{cov(\Delta ln(c_{i,t}), \eta_{i,t})}{var(\eta_{i,t})} &= \underbrace{E[\phi^{\eta}_{i,t}]}_{\phi^{\eta}}.
\end{aligned}$$

Another way to obtain such expressions is to relax the assumption that shocks are normal but take a first order approximation of log-consumption growth around small income shocks and neglect all higher-order terms. This would mean imposing that all higher order moments of the distribution are zero instead of imposing that the odd order moments are zero and the even order moments multiples of the second order moment.

2.2 Identification strategy: instrumenting with future income growth

When the realizations of the shocks ε and η are observed, typically in the context of natural experiments, it is straightforward to measure the covariances $cov(\Delta ln(c_{i,t}), \varepsilon_{i,t})$, $cov(\Delta ln(c_{i,t}), \eta_{i,t})$, and variances $var(\varepsilon_{i,t})$, $var(\eta_{i,t})$, thus to measure the elasticities. An OLS estimator of the effect of each shock on log-consumption growth yields unbiased estimates of the elasticities above. In survey data, however, the realizations of the shocks ε and η are not directly accessible. Only c and y are reported. The objective of the identification strategy is then to express the unobserved covariances and variances of the income shocks from observed moments of log-income and log-consumption growth.

The solution I use is to instrument the effect of current income growth on consumption growth with future income growth. Indeed, future income growth correlates with the realization of the current transitory shock but not with the realization of the current permanent shock. A transitory income shock raises income at t, inducing a positive income growth between $t-1$ and t, but as its effect dissipates it does not raise income by as much at $t+1$ and even less so at $t+2$, inducing a negative income growth between t and $t+1$ and between $t+1$ and $t+2$. On the contrary, a permanent shock increases income once and for all, so it generates a positive income growth between $t-1$ and t and has no effect on income growth between t and subsequent periods. The fluctuations in current log-consumption growth that are driven by the transitory shock can be identified as the part of current log-consumption growth that covaries with future log-income growth (and similarly the fluctuations in current log-income growth that are driven by the transitory shocks correspond to the part of log-income growth that covaries with future log-income growth):

$$
\begin{aligned}
cov(\Delta ln(c_{i,t}), -\Delta ln(y_{i,t+2})) &= \theta cov(\Delta ln(c_{i,t}), \varepsilon_{i,t}), \\
cov(\Delta ln(y_{i,t}), -\Delta ln(y_{i,t+2})) &= \theta var(\varepsilon_{i,t})
\end{aligned}
$$

An unbiased estimator of the elasticity of consumption to a transitory shock, robust to the presence of a correlation between consumption growth and past shocks is:

$$
\hat{\phi}^{\varepsilon} = \frac{cov(\Delta ln(c_{i,t}), \Delta ln(y_{i,t+2}))}{cov(\Delta ln(y_{i,t}), \Delta ln(y_{i,t+2}))} = \phi^{\varepsilon}
$$

2.3 Bias from correlation with past shocks in non-robust estimators

Neglecting the presence of a correlation between consumption growth and the realizations of past income shocks leads a researcher to use additional moments in which they overlook some terms, inducing a bias.

2.3.1 Semi-structural estimators

Elasticity to a transitory shock (BPP) When past income shocks have not effect on current log-consumption growth, it is possible to use both future income growth between t and $t+1$ and future income growth between $t+1$ and $t+2$ as instruments, to get more identifying restrictions. The covariance between log-consumption growth and the transitory shock is overidentified because it is measured from two restrictions. Yet, if in reality past income shocks do affect log-consumption growth, using future income growth between t and $t+1$ as an instrument and neglecting the correlation between log-consumption growth and past income shocks leads to a bias:

$$cov(\Delta ln(c_t), -\Delta ln(y_{t+2})) = \theta cov(\Delta ln(c_t), \varepsilon_t),$$
$$cov(\Delta ln(c_t), -\Delta ln(y_{t+1})) = (1-\theta) cov(\Delta ln(c_t), \varepsilon_t) + \theta \underbrace{cov(\Delta ln(c_t), \varepsilon_{t-1})}_{<0 \text{ (neglected)}}.$$

The variance of the transitory income shock is estimated as before:

$$cov_t(\Delta ln(y_t), -\Delta ln(y_{t+2})) = \theta var_t(\varepsilon_t).$$

Thus, the consumption elasticity to a transitory shock is identified from the two following expressions:

$$\hat{\phi}^{\varepsilon}_{BPP} = \frac{cov(\Delta ln(c_t), -\Delta ln(y_{t+2}))}{cov(\Delta ln(y_t), -\Delta ln(y_{t+2}))} = \phi^{\varepsilon},$$
$$\hat{\phi}^{\varepsilon}_{BPP} = \frac{\theta}{1-\theta} \frac{cov(\Delta ln(c_t), -\Delta ln(y_{t+1}))}{cov_t(\Delta ln(y_t), -\Delta ln(y_{t+2}))} = \phi^{\varepsilon} + \frac{\theta}{1-\theta} \underbrace{\frac{cov(\Delta ln(c_t), \varepsilon_{t-1})}{var_t(\varepsilon_t)}}_{<0 \text{ (neglected)}} < \phi^{\varepsilon},$$

The first expression is the same as that of the robust estimator and is unbiased. The second expression neglects the fact that a high realization of the past transitory shock at $t-1$ reduces log-income growth at $t+1$, raising $-\Delta ln(y_{t+1})$, and reduces log-consumption consumption at t for precautionary reason, inducing a negative covariance between the two. It induces a downward bias in the estimation: the negative impact of the past transitory shock on the covariance is erroneously attributed to the fact that the covariance between log-consumption growth and the current transitory shock is smaller than than it actually is. The bias is akin to a problem of instrument endogeneity: the variable $\Delta ln(y_{t+1})$ is used as an instrument, to filter the fluctuations in $\Delta ln(y_t)$ that are driven by a transitory shock and get rid of the fluctuations driven by a permanent shock. Yet because this instrument covaries with the dependent variable $\ln(c_t)$ through both the current transitory shock and the past transitory shocks, it is not exogenous: there is a bias.

Erroneously assuming an MA(0) transitory income process when the true process is an MA(1) leads to an even more severe bias. It induces a researcher to erroneously use the moment:

$$cov(\Delta ln(c_t), -\Delta ln(y_{t+1})) = cov(\Delta ln(c_t), \varepsilon_t) \underbrace{-\theta cov(\Delta ln(c_t), \varepsilon_t)}_{<0 \text{ (neglected)}} + \underbrace{\theta cov(\Delta ln(c_t), \varepsilon_{t-1})}_{<0 \text{ (neglected)}} < \phi^{\varepsilon},$$

neglecting both the terms $-\theta cov(\Delta ln(c_t), \varepsilon_t) < 0$ and $\theta cov(\Delta ln(c_t), \varepsilon_{t-1}) < 0$ that appear in the true moment. If the true process for transitory income is an MA(0), however, this method yields an unbiased estimator. Indeed, when $\theta = 0$, there is only one instrument available, future income growth between t and $t+1$, and it gives an unbiased estimate.

Elasticity to a permanent shock (BPP) The problems caused by the presence of a correlation with past shocks are more serious for the identification of the elasticity of consumption to a permanent shock: I did not manage to find restrictions that would build a consistent estimator in the presence of a correlation between log-consumption growth and past shocks. When past shocks have no effect on log-consumption growth, it is possible to identify the elasticity using the sum of current and future log-income growth as an instrument. Indeed, because a transitory shock only increase log-income for two periods, it has no impact on log-income growth between $t-1$ and $t+2$. On the contrary, a permanent shock at t raises log-income permanently thus raises log-income growth between $t-1$ and $t+2$. The covariance between log-consumption growth and the permanent shock at t is then identified as the part of the fluctuations in log-consumption growth or log-income growth that covaries with log-income growth between $t-1$ and $t+2$:

$$cov(\Delta ln(c_t),\ \Delta ln(y_t)+\Delta ln(y_{t+1})+\Delta ln(y_{t+2})) = cov(\Delta ln(c_t),\eta_t)+\underbrace{cov(\Delta ln(c_t),-\varepsilon_{t-1}-\theta\varepsilon_{t-2})}_{>0\text{ (neglected)}}.$$

In the case of the identification of the variance, by construction, log-income growth depends on the realizations of past transitory shocks up to $t-2$. It is necessary to have an instrument that covaries with the current permanent shock but not with neither the current transitory shock nor past transitory shocks up to $t-2$. Log-income growth between $t-3$ and $t+2$ verifies these conditions: transitory shocks at $t-2$, $t-1$ and t have no impact on its value, while the permanent shock at t covaries positively with it. The variance is identified with:

$$cov(\Delta ln(y_t),\Delta ln(y_{t-2})+\Delta ln(y_{t-1})+\Delta ln(y_t)+\Delta ln(y_{t+1})+\Delta ln(y_{t+2})) = var(\eta_t).$$

The estimator of the elasticity to permanent shocks is:

$$\hat{\phi}^{\eta}_{BPP} = \frac{cov(\Delta ln(c_t),\ \Delta ln(y_t)+\Delta ln(y_{t+1})+\Delta ln(y_{t+2}))}{var_t(\eta_t)} = \phi^{\eta}+\underbrace{\frac{cov(\Delta ln(c_t),-\varepsilon_{t-1}-\theta\varepsilon_{t-2})}{var(\eta_t)}}_{>0\text{ (neglected)}} > \phi^{\eta}.$$

In the presence of a correlation between log-consumption growth and past income shocks, this estimator is biased upwards. Past transitory shocks correlates negatively with log-income growth between $t-1$ and $t+2$, and negatively with log-consumption growth between $t-1$ and t, raising the observed covariance between the instrument and log-consumption growth. When their effect is neglected, the increase in the covariance is attributed to the fact that consumption responds more to a permanent shock than it actually does.

Erroneously assuming that transitory income is an MA(0) process leads to a bias as well. If the true process for transitory income is an MA(0) and not an MA(1), it remains impossible for me to find an instrument that would identify the elasticity of consumption to a permanent shock (all instruments that capture the effect of the permanent shock also depend on past income shocks).

Elasticity to a permanent shock (Kaplan and Violante (2010)) Kaplan and Violante (2010) identify the elasticity of consumption to a permanent shock using a different instrument than the original one used in BPP. They model transitory income as an MA(0) process but rely on log-income growth between $t-2$ and $t+1$ to instrument fluctuations in log-consumption growth instead of using only log-income growth between $t-1$ and $t+1$:[9]

$$cov(\Delta ln(c_t), \Delta ln(y_{t-1}) + \Delta ln(y_t) + \Delta ln(y_{t+1})) = cov(\Delta ln(c_t), \eta_t) + \underbrace{cov(\Delta ln(c_t), \eta_{t-1} - \varepsilon_{t-2})}_{\text{(neglected)}}.$$

The variance is estimated with the same instrument, which is the same in BPP but adjusted for the fact that transitory income is an MA(0) process:

$$cov(\Delta ln(y_t), \Delta ln(y_{t-1}) + \Delta ln(y_t) + \Delta ln(y_{t+1})) = var(\eta_t).$$

The Kaplan and Violante (2010) version of the BPP estimator of the elasticity to a permanent shock is:

$$\hat{\phi}^{\eta}_{BPP-KV} = \frac{cov(\Delta ln(c_t),\ \Delta ln(y_{t-1}) + \Delta ln(y_t) + \Delta ln(y_{t+1}))}{var(\eta_t)} = \phi^{\eta} + \underbrace{\frac{cov(\Delta ln(c_t), \eta_{t-1} - \varepsilon_{t-2})}{var_t(\eta_t)}}_{\text{(neglected)}} \neq \phi^{\eta}.$$

This identification method is biased as well, but this time the direction of the bias is undetermined. Indeed, the instrument, log-income growth between $t-2$ and $t+1$, correlates with log-consumption growth through both past transitory and past permanent shocks. I proved that the realizations of past transitory shocks are strictly negatively correlated with current log-consumption growth but I do not know in which direction goes the effect of the realizations of past permanent shocks (having experienced a good permanent shock in the past increases the expected amount of resources of a household but also increases the variance of future income possibly strengthening the need for precautionary saving). Thus, the instrument is endogenous and the estimator biased but the direction of the bias is undetermined. Note that Kaplan and Violante (2010) find that their version of the BPP estimator underestimates the true elasticity to permanent shocks: it indicates that, for their particular calibration at least, either past permanent shocks reduce log-consumption growth, as transitory shocks do, or past permanent shocks raise log-consumption growth a little but not so much that their effect offsets that of the past transitory shock at $t-2$.[10]

2.3.2 Structural estimators

Later work have built a different estimation strategy that builds on the full life-cycle model rather than on some of the restrictions it implies.[11] The method consists in deriving approximated analytical expressions for the consumption elasticities to shocks, and estimating their components from survey data. In the standard life-cycle model presented above, the true elasticities coincide with the percentage change

[9] In the theoretical description of their identification (if transitory income is MA(0)), BPP presents the covariance between log-consumption growth and log-income growth between $t-2$ and $t+1$ as their identifying moment but in practice the moment they use is the covariance between log-consumption growth and log-income growth between $t-1$ and $t+1$ because they set the covariance between log-consumption growth and log-income growth between $t-2$ and $t-1$ to zero.

[10] Because the two elasticities are estimated separately and not jointly in Kaplan and Violante (2010), this bias in the measure of the elasticity to permanent shocks cannot be partly attributed to the fact that the bias affecting the measure of the elasticity to transitory shock is contaminating their estimate, as could be the case in the original BPP estimator.

[11] See for example Blundell, Pistaferri, and Saporta-Eksten (2016), Pistaferri, Saporta-Eksten, and Blundell (2017), and Blundell, Graber, and Mogstad (2016)

in total expected resources net of total expected precautionary consumption growth:

$$\phi^{\varepsilon} = \frac{\left(\frac{dW_{t+1}}{d\varepsilon_{t+1}}\right)^0 - \overbrace{\left(\frac{dPG_{t+1}}{d\varepsilon_{t+1}}\right)^0}^{\text{precaution (1)¡0}}}{W_{t+1} - \underbrace{(PG_{t+1})^0}_{\text{precaution (2)¿0}}} \quad \text{and} \quad \phi^{\eta} = \frac{\left(\frac{dW_{t+1}}{d\eta_{t+1}}\right)^0 - \overbrace{\left(\frac{dPG_{t+1}}{d\eta_{t+1}}\right)^0}^{\text{precaution (1)} \neq 0}}{W_{t+1} - \underbrace{(PG_{t+1})^0}_{\text{precaution (2)¿0}}}.$$

with:

$$\begin{cases} W_{t+1} = (1+r)a_{t+1} + \sum_{s=0}^{T-t-1} \frac{E_{t+1}[y_{t+1+s}]}{(1+r)^s} & \text{the amount of total expected resources at } t+1 \\ PG_{t+1} = \sum_{s=1}^{T-t-1} l_{t+1,s} \frac{E_{t+1}[\varphi_{t+s} + \tilde{\lambda}_{t+s}]}{(1+r)^s} & \text{the total expected precautionary growth at } t+1 > 0 \end{cases}$$

Therefore, precautionary behavior affects the elasticity of consumption in two ways, indicated with (1) and (2). First, denoted with (1), an income shock modifies the amount of precautionary saving that a household needs to make, which can attenuate of reinforce their consumption response: if the shock increases income but raises precautionary saving, consumption responds less than it would under perfect foresight, while if the shock increases income and reduces precautionary saving, consumption responds more. Second, denoted with (2), because a household saves for precautionary reasons, it consumes less at $t+1$ than it would under perfect foresight. As a result, a given change in consumption corresponds to a larger percentage change in consumption for a household engaging in precautionary behavior than for a household with perfect foresight: the denominator of its elasticity of consumption is always smaller.

Because they make the same approximation as BPP and neglect the precautionary component of consumption growth, existing semi-structural estimators obtain the following expression:

$$(\phi_t^{\varepsilon})^{BPS} = \frac{\left(\frac{dW_{t+1}}{d\varepsilon_{t+1}}\right)^0}{W_{t+1}^0} = (\phi_t^{\varepsilon})^{PF} < \phi^{\varepsilon} \quad \text{and} \quad (\phi_t^{\eta})^{BPS} = \frac{\left(\frac{dW_{t+1}}{d\eta_{t+1}}\right)^0}{W_{t+1}^0} = (\phi_t^{\eta})^{PF} \neq \phi^{\eta}.$$

It coincides with the elasticity that is obtained in the absence of uncertainty, or under quadratic preferences. The resulting estimate of the elasticity of consumption to a transitory shock is downward biased because they are neglecting effects (1) and (2).

3 Empirical Implementation

3.1 Data

PSID 1978-1992: My main dataset is the same as in BPP. It contains observations from the Panel Study of Income Dynamics (PSID) between 1978 and 1992.[12] For comparison purpose, I apply the same selection as BPP, which is detailed in their paper. The final sample is composed of 15,779 observations from 1,765 households[13].

[12] I considered including additional years after 1992, and before the 1999 remodeling of the survey, but a number of the questions used by Blundell, Pistaferri and Preston are also redesigned in 1993, in particular regarding financial income, and the impact of these changes is difficult to measure.

[13] My sample is exactly the same as that of BPP. The number of observations reported differ (they report 17,604 observations) simply because they count observations of log-income, while I report observations of log-income growth. Among these 15,779 observations, 12,041 comprise an observation of both log-income growth and log-consumption growth (and demographics),

The variables are constructed as follows. Net income is the taxable family income reported by the household, net of income from financial assets and then net of the federal taxes paid on nonfinancial income, deflated by the Consumer Price Index (CPI)[14]. Gross earnings is the taxable family income, net of financial assets and deflated by the CPI. The real labor income of the male head is the yearly revenue he declares, deflated by the CPI. His wage rate is his real labor income divided by the hours he worked. The yearly number of hours worked is obtained from multiplying the average number of hours worked per week by the number of weeks worked. Questions on income are retrospective and refer to the previous calendar year.

The PSID only reports consumption expenditures on food, while it is more interesting to use a broader category of nondurable consumption for the present exercise, in particular because food expenditures are known to be less elastic to income shocks than expenditures on broader categories of consumption. To overcome the problem, nondurable consumption is imputed from demographics and food consumption, with the coefficients used for the imputation computed from the Consumer Expenditure Survey (CEX) over the same period. Further details are provided in the original BPP paper. Nondurable consumption is the sum of expenditures on food (at home and away from home), alcohol, tobacco, nondurable services, heating fuel, public and private transport (including gasoline), personal care, clothing and footwear. In particular, this definition excludes expenditure on housing, health, and education. To obtain the real analog to nominal consumption, it is deflated by the CPI.

I detrend variables from the impact of the demographic variables that can shift log-income and the utility derived from consumption, interacted with the period. To do so, I regress log-income and log-consumption on year and year-of-birth dummies, and on dummies for education, race, family size, number of children, region, employment status, residence in a large city, and outside dependent, interacted with a period dummy. In the robustness section, I test the sensitivity of my results to the variables included in this set of demographics.

3.2 Estimating with a generalized method of moment

I estimate the elasticity with a generalized method of moments. The statistical model implies that:

$$cov(\Delta ln(c_{i,t}), \Delta ln(y_{i,t+2})) = E[\Delta ln(c_{i,t}) \Delta ln(y_{i,t+2})] = \phi^{\varepsilon} E[\varepsilon_{i,t}^2]$$
$$cov(\Delta ln(y_{i,t}), \Delta ln(y_{i,t+2})) = E[\Delta ln(y_{i,t}) \Delta ln(y_{i,t+2})] = E[\varepsilon_{i,t}^2]$$

Thus, the following moment restrictions holds:

$$E[\underbrace{(\Delta ln(c_{i,t}) - \varphi^{\varepsilon} \Delta ln(y_{i,t})) \Delta ln(y_{i,t+2})}_{g(X_{i,t}, \phi^{\varepsilon})}] = 0$$

and 3,738 comprise only log-income growth (and demographics).

[14]Federal taxes on nonfinancial income are assumed to be a proportion of total federal taxes; the proportionality coefficient is given by the ratio between nonfinancial income and total income.

with $X_{i,t} = (\Delta ln(c_{i,t}), \Delta ln(y_{i,t}), \Delta ln(y_{i,t+2}))$ and $\phi^{\varepsilon} = E[\phi^{\varepsilon}_{i,t-1}]$. The elasticity ϕ^{ε} is computed as the value that minimizes a norm of the sample analog of this moment:

$$\hat{\phi^{\varepsilon}} = \underset{\phi^{\varepsilon}}{\text{argmin}} \left(\frac{1}{L}\frac{1}{I}\sum_{t=1}^{L}\sum_{i=1}^{I} g(X_{i,t}, \phi^{\varepsilon}) \right)^{\mathsf{T}} \hat{W} \left(\frac{1}{T}\frac{1}{I}\sum_{t=1}^{T}\sum_{i=1}^{I} g(X_{i,t}, \phi^{\varepsilon}) \right)$$

with L the number of periods observed, I the number of households of the sample, and $\hat{W}$ the weighting matrix. The weighting matrix is chosen to be robust to heteroskedasticity. The residual are clustered by year $\times$ education level. This means that I allow for the shocks received at a given period by two households with the same education level to be correlated (e.g. not only are they possibly drawn from different distributions but the realization of one can help predict the realization of the other); the shocks received in the same year by households with different education levels or the shocks received in different years by households with the education level as well, however, are assumed to be uncorrelated.

4 Results

4.1 Covariances between log-consumption growth and log-income growth

Table 1: Covariances between $\Delta ln(c)$ and past, present and future $\Delta ln(y)$

Covariances	$\Delta ln(y_{i,t-1})$	$\Delta ln(y_{i,t})$	$\Delta ln(y_{i,t+1})$	$\Delta ln(y_{i,t+2})$	$\Delta ln(y_{i,t+3})$
$cov(\Delta ln(c_{i,t}), .)$	-0.0028^{*}	0.0139^{***}	-0.0017	-0.0028^{*}	0.0012
	(0.0016)	(0.0017)	(0.0016)	(0.0015)	(0.0014)
$cov(\Delta ln(y_{i,t}), .)$	-0.0247^{***}	0.0816^{***}	-0.0278^{***}	-0.0059^{***}	0.0013
	(0.0017)	(0.0026)	(0.0016)	(0.0013)	(0.0012)
Obs.	8,112	8,958	8,958	8,958	7,819

Note: This table presents the covariances between the log-consumption growth and log-income growth of a household and its past, present or future log-income growth among household-year observations (i,t) for which $\Delta ln(c_{i,t})$, $\Delta ln(y_{i,t})$, and $\Delta ln(y_{i,t+2})$ are observed, which are the household-periods present in the estimation sample.

Before looking at the elasticity estimate, Table 1 presents the value of empirical moments, to verify whether they are consistent with the statistical model. Each covariance is computed unconditionally over the household-year observations for which it exists, among the household-year observations that are used for estimation, that is, those for which $\Delta ln(c_{i,t})$, $\Delta ln(y_{i,t})$, and $\Delta ln(y_{i,t+2})$ are observed. The covariances of log-consumption are consistent with my general statistical model, but incompatible with a model that does not allow for a correlation between log-consumption growth and past shocks for two reasons. First, such a model predicts that the correlation between log-consumption growth and past log-income growth is zero. In the data, this covariance is large (its magnitude is bigger than the correlation between log-consumption growth and future log-income growth), and statistically different from zero at the 10% level. In my general statistical model, the covariance with past log-income growth captures the correlation between the realizations of past income shocks and the current consumption growth of a household, which can be non-zero. Second, a model that does not allow for a correlation with past shocks predicts that the covariance between log-consumption growth and future log-income growth at $t+1$ is larger in absolute value than its covariance with log-income growth at $t+2$ by a factor of $\frac{1-\theta}{\theta}$, which is approximately 8 with BPP's estimate of $\theta = 0.11$. In the data, the opposite is true, and the covariance between log-consumption growth and log-income growth at $t+1$ is smaller than its covariance with log-income growth at $t+2$. With a statistical model that allows for a correlation between consumption growth and past income shocks, however, the covariance with log-income growth at $t+1$ is not necessarily larger, because log-income growth at $t+1$ partly incorporates the realizations of past transitory shocks that can covary negatively with current log-consumption growth and lessen the magnitude of the covariance between log-consumption growth and log-income growth at $t+1$, bringing its value it close to or below that of the covariance with log-income growth at $t+2$.

The covariances of log-income growth are consistent with a transitory-permanent income process in which the transitory component of income is an MA(1). If the transitory component was an MA(0), the covariance between log-income growth at t and at $t+2$ should not be statistically different from zero, while it is. If the transitory component was an MA(2), the covariance between log-income growth at t and at $t+3$ should be statistically different from zero, while it is not. If permanent income was not a random walk but an AR(1) with a coefficient different from one, it means that the effect of a permanent shock would still affect the permanent income of a household until the last period of the its life but with its effect fading away at each period by a factor ρ. In that case and if the effect was large, the autocovariance between log-income growth at t and at $t+3$ should be statistically different from zero, while it is not.

4.2 Average elasticity of nondurable consumption to a transitory shock

Table 2: Estimated elasticities ϕ - Baseline

	Robust	Non-robust	BPP (non-robust)
ϕ^{ε}	0.465***	0.071**	0.053
	(0.172)	(0.032)	(0.043)
Implied MPC	0.388***	0.059**	0.041
	(0.143)	(0.027)	(0.034)
Obs.	8,958	8,958	12,041
Moments used	(1), (2)	(1), (2), (3)	$(1)_t$, $(2)_t$, $(3)_t$ $\forall t$
		$\theta = 0.11$	+ others
Clustering	year × educ	year × educ	no

Note: This table presents estimates of the average elasticity of nondurable consumption expenditures to shocks on net income. The estimator is a GMM with a robust weighting matrix. When clustered, it is so by year × education level. When the elasticity cannot be identified separately from θ, I provide an external value for θ that is the one measured by BPP: $\theta = 0.11$. The main estimating moments are:

$$0 = cov(\Delta ln(y_{i,t}), \Delta ln(y_{i,t+2})) - \theta var(\varepsilon_{i,t}) \quad (1)$$
$$0 = cov(\Delta ln(c_{i,t}), \Delta ln(y_{i,t+2})) - \phi^{\varepsilon} \theta var(\varepsilon_{i,t}) \quad (2)$$
$$\underbrace{0 = cov(\Delta ln(c_{i,t}), \Delta ln(y_{i,t+1})) - \phi^{\varepsilon}(1-\theta)\theta^{-1} var(\varepsilon_{i,t}))}_{\text{does not hold if correlation with past shocks because of a missing term}} \quad (3)$$

The first column on the left presents the results from a specification that is robust to a correlation with past shocks because it uses only a weighted sum of moments (1) and (2); the second column presents the results obtained when using a specification that is biased in the presence of a correlation with past shocks because it uses a weighted sum of moments (1) and (2) but also of moments (1) and (3); the third column presents the results obtained with the original BPP estimator that uses moments (1), (2), and (3) conditionally on the period t at which log-consumption and log-income growth are observed, and also uses additional estimating moments. A household-year pair is counted as observed if at least one of the estimating moments involving the coefficient ϕ^{ε} is observed for this pair.

Robust estimator The first column of Table 2 reports the results obtained with my robust estimator, which is based on a combination of moments that hold even in the presence of a correlation between log-consumption growth and past income shocks. It estimates the elasticity of consumption to a transitory income shock to be large, with a point estimate of 0.465. This means that on average in the sample a 10% increase in transitory income—which raises current income by 10% and next period income by $(\theta * 100)$%—leads to a 4.6% increase in current consumption. The estimate is statistically different from zero at the 1% level.

Non-robust estimators The two other columns present the results obtained with estimators that are not robust to the presence of a correlation between log-consumption growth and past income shocks, because they rely on moment (3), an expression of the covariance between log-consumption growth and log-income growth at $t+1$ that neglects a non-zero term in the presence of a correlation. These two estimators yield much smaller estimates of the elasticity of consumption to a transitory income shock than the robust estimator. The second column features the results from an estimator that uses the same moment as the robust estimator plus moment (3), the estimation method being otherwise identical to that of the first column. The associated estimate of the elasticity of consumption to a transitory income shocks is 0.037, much below this estimate obtained with a robust estimator. The estimate is precisely measured and statistically different from zero at the 5% level. The fact that the point-estimate is much smaller is in line with my theoretical prediction that, in the presence of a negative correlation between log-consumption growth and past transitory income shocks, using moment (3) as an estimating moment induces a downward bias in the measure of the elasticity of consumption to a transitory income shock. The third column presents the results obtained with the original BPP estimator, which uses moment (3) as well, but differ from the robust estimator on other grounds because it uses other additional moments, uses all moments conditionally on the period, and has one additional detrending variable.[15] The point estimate of the elasticity of consumption to a transitory income shock is similar to that of the other non-robust estimator, at 0.053, and not statistically different from zero. This suggests that the additional estimating moments used slightly reinforce the downward bias associated with the moment (3), and decrease the precision of the estimation.

Marginal Propensity to Consume (MPC) Structural methods measure the elasticity of consumption to a transitory income shock, that is, the percentage change in consumption associated with a percentage change in income. On the contrary, natural experiments typically measure the MPC out of transitory income, that is, the level change in consumption associated with a level change in income. To facilitate the comparison with natural experiments, I derive the MPC out of transitory income that is implied my estimated value of the consumption elasticity. To move from one measure to another, I proceed as follows. In level, net income at t writes $y_t = e^{p_t} e^{\varepsilon_t} e^{\theta \varepsilon_{t-1}}$. A good transitory income shock $d\varepsilon_t$ generates a transitory income gain $dy_t = d\varepsilon_t y_t$. Natural experiments measure the impact of such a gain on consumption: $MPC = \frac{dc_t}{dy_t} = \frac{dc_t}{d\varepsilon_t}\frac{1}{y_t}$. In comparison, the BPP estimator measures the elasticity to transitory shocks, $\phi^\varepsilon = \frac{dln(c_t)}{d\varepsilon_t} = \frac{dc_t}{d\varepsilon_t}\frac{1}{c_t}$. Therefore, to derive the MPC from the elasticity, I compute:

$$MPC = \frac{dc_t}{d\varepsilon_t}\frac{1}{y_t} = \frac{dc_t}{d\varepsilon_t}\frac{c_t}{c_t}\frac{1}{y_t} = \phi^\varepsilon \frac{c_t}{y_t}.$$

The result is presented in the second line of Table 2. In the sample, the average ratio $\frac{c_t}{y_t}$ is 0.752, so MPC implied by an elasticity of 0.465 is $0.752 \times 0.465 = 0.388$. This means that 39% of a transitory income gain is consumed over the following year.

Comparison with the literature on natural experiments How does this value of 39% compare with the results derived from natural experiments? Over a similar period, between 1979 and 1990, Souleles (1999) exploits tax refunds to measure the MPC out of transitory income. He estimates that 9% of a tax refund is spent on strictly nondurable expenditures over the quarter following receipt, and his estimate is statistically significant at the 5% level (Table 5 in his paper). His definition of strictly nondurable consumption is a little more restrictive than the one used in this paper and in BPP, so that his 9% estimate

[15] I discuss the effect of these additional features in Appendix C, and I check the effect of the additional detrending variable in the robustness section.

is a lower bound for the MPC associated with the more general category of nondurable consumption goods considered here. Papers that investigate the impact of tax rebates obtain larger estimates of the marginal propensity to consume nondurable goods: studies converge to a quarterly value of 0.25.[16] The effect of the transitory shock on spending in the second three-month period after the payment is not usually statistically significant, so the yearly MPC is larger but the share consumed over the next quarters is usually not precisely estimated. Yet, these tax rebates have a number of specificities that may differentiate them from the average transitory shock captured in survey data. First, they typically occur during an economic recession, which are periods when consumers are on average facing more uncertainty and more likely to have low levels of income and assets. As consumers with lower levels of assets and income can be responding more to an income shock, for a number of reasons, the period of a tax rebate would be one in which consumers respond more to a transitory shock than they do on average over a long period. Second, both tax refunds and rebates are likely to be more or less anticipated—consumers should have some idea of whether they have been overpaying or not and- fiscal stimulus have been announced. This has the opposite effect: consumers would respond less to a tax rebate than to the average transitory shock they face. Third, a tax rebate may be accompanied by an unobserved change in households' expectations: on the one hand if households simultaneously anticipate an increase in future taxes, they might respond less to a tax rebate than to a typical transitory shock; on the other hand, if they have some belief that the rebates are going to have some persistence, than might respond more than to a typical, clearly non-persistent or shortly persistent income shock. Using data from lottery gains in Norway,Fagereng, Holm, and Natvik (2016) obtain a very precisely estimated yearly MPC out of a lottery gain, and their point estimate is 0.34. To sum it up, two characteristics seem robust to the idiosyncrasies of the transitory income variations considered in different natural experiments: (i) the MPC out of an unexpected transitory gain is statistically significant (ii) its value over the year following the gain is above 9%. The MPC of 39% derived from my generalized estimator is consistent with both stylized facts, while the MPC derived from the BPP estimator conflicts with both.

[16]Johnson, Parker, and Souleles (2006) study the response of consumption to the 2001 fiscal stimulus implemented in the U.S. They obtain that the consumption of nondurable goods increased by 0.38 of the rebate, within quarter following receipt. Kaplan and Violante (2014) do a similar correction and obtain close results. Misra and Surico (2014) refine the technique to account for heterogeneity in the response of consumption and obtain a marginal propensity to spend on nondurable goods of 0.25. Similar findings are obtained for the 2008 tax rebate (Parker, Souleles, Johnson, and McClelland (2013), Misra and Surico (2014).

4.3 Differences in elasticity across goods

Table 3: Estimated elasticities ϕ - Comparison between different sets of consumption goods

	Nondurable	Food	Total
ϕ^{ε}	0.465***	0.242*	0.527**
	(0.219)	(0.139)	(0.219)
Implied MPC	0.388***	0.042**	1.083**
	(0.143)	(0.024)	(0.448)
Obs.	8,958	8,973	8,958

Note: This table presents estimates of the average elasticity of consumption expenditures to shocks on net income for different categories of consumption expenditures. The firs column on the left shows the elasticities associated with nondurable expenditures, the second column with food expenditures, and the last column on the right with total expenditures (nondurables and durables).

Table 5 presents the elasticities obtained when measuring the response of two other categories of consumption expenditures: food expenditures and total consumption expenditures, which includes both nondurable and durable expenditures. The elasticity of food expenditures is 0.242 and is statistically significant at the 10% level. This smaller point estimate suggests that food expenditures are less elastic than other nondurable expenditures. The associated MPC is 3%. Such a small number is not surprising: if food expenditures represent a small share of the expenditures of the household, the percentage of an income gain that is devoted to food is mechanically small. The elasticity to total expenditures is 0.527, significant at the 5% level, and the MPC is 108%, suggesting that on average all of the transitory income gain is consumed in the following year. To compare those result with the literature on natural experiments, Souleles (1999) finds that the average quarterly MPC associated with food expenditures is 6% and the average quarterly MPC associated with total expenditures is 64%, both being statistically significant. Parker, Souleles, Johnson, and McClelland (2013) obtains that the average quarterly MPC associated with food is around 2%, and the average quarterly MPC associated with total consumption is 54%. Both sets of results are consistent with my yearly estimates, except for my MPC of food consumption that is a little small in comparison with the measure of Souleles.

These results are useful on two other regards. First, because food consumption is not imputed but is reported together with income in the PSID, this serves as a test for the validity of the imputation procedure. Results are favorable, since the pattern is qualitatively similar with food consumption as with imputed consumption: with a robust estimator, the elasticity of food expenditures to a transitory shock is large (although smaller than the elasticity of all nondurable expenditures) and statistically different from zero at the 10% level.

Second, in a statistical model such that equation (3.1) holds for the three categories of consumption expenditures at the same time (e.g. such that I am not changing model when changing the categories of consumption expenditures)[17], the comparison between the elasticities provide information about the

[17] It is the case that equation (3.1) holds for the three categories of consumption expenditures when a household solving a

impact of a transitory income shock on the composition of the consumption basket. The elasticity of the share of food expenditures c_t^{Food} in all nondurable expenditures c_t is:

$$\frac{1}{\frac{c_t^{Food}}{c_t}} \frac{d\frac{c_t^{Food}}{c_t}}{d\varepsilon_t} = \phi^{\varepsilon Food} - \phi^{\varepsilon} = -0.22,$$

with $\phi^{\varepsilon Food}$ and ϕ^{ε} are the elasticities of food expenditures and nondurable expenditures to a transitory income shock. This means that, following a transitory income gain that increases current income by 10% of its initial value, the share of food expenditures in nondurable expenditures decreases by 2.2%. Similarly, the elasticity of the share of nondurables expenditures c_t in total expenditures c_t^{Tot} is:

$$\frac{1}{\frac{c_t}{c_t^{Tot}}} \frac{d\frac{c_t}{c_t^{Tot}}}{d\varepsilon_t} = \phi^{\varepsilon} - \phi^{\varepsilon Tot} = -0.06,$$

with ϕ^{ε} and $\phi^{\varepsilon Tot}$ the elasticities of nondurable expenditures and total expenditures to a transitory income shock. It implies that, following a 10% transitory increase in current income, the share of nondurable expenditures in total expenditures decreases by 0.6%.

4.4 Differences in elasticity across households

Table 4: Estimated elasticities ϕ - Comparison between different subgroups of households

	Financial income ($\leq/>$ 300$ per year)		Wage rate ($</\geq$ 13$ per hour)		Age ($\leq/>$ 44 years old)	
	Low	High	Low	High	Low	High
ϕ^{ε}	0.542	0.398*	0.488**	0.311	0.692*	0.318
	(0.369)	(0.236)	(0.221)	(0.856)	(0.488)	(0.225)
Obs.	4,556	4,402	5,312	3,646	5,030	3,928

Note: This table presents estimates of the average elasticity of nondurable consumption expenditures to shocks on net income within different subgroups of the population.

Table 3 presents the estimates obtained when implementing the estimator separately on different subcategories of the population. Note that the variables are still detrended over the whole sample and not over their subsample, to keep the effect of demographic characteristics identical across subsamples and facilitate the comparison. The main result from this table is that the point estimate of the elasticity to a transitory shock is larger for households with lower financial income, a lower wage rate of the head, and a younger head, but is not close to zero for households with higher financial income, a higher wage rate, and an older head and the difference across categories are not statistically significant. More precisely, households with low financial income have an average elasticity to transitory shocks of 0.542. This is larger than the elasticity of 0.398 that is estimated among households with high financial income. Yet for both categories the elasticity is large, and it is statistically significant at the 10% level for household with high financial income. The higher precision of the estimation could come from the fact that households

life-cycle model derives utility separately from these three categories of consumption expenditures.

with high financial income are subject to more transitory income variations. The results are qualitatively similar for households whose head has a low versus a high wage rate: the elasticity of consumption is larger among households with a low wage rate than among households with a high wage rate, at 0.488 and 0.311, but both elasticities are large. A household whose male head is below 44 responds on average twice as much to a transitory income shock than a household whose head is strictly above 44. A possible explanation is that young households face on average more uncertainty and are more often constrained.

These results on the heterogeneity of the consumption response across categories match with the findings from natural experiments. Parker, Souleles, Johnson, and McClelland (2013) document the impact of the 2007/2008 tax rebate across households with different characteristics. They find no statistically significant impact on the marginal propensity to consume of being in low versus a high group of age, income, liquid assets or housing status.[18] Fagereng, Holm, and Natvik (2016) obtain that all categories of the population respond significantly to a lottery gain.

Finally, in the three cases considered, the weighted average of the elasticities of each subcategory coincides quite closely with the elasticity of the whole population. This is consistent with my result that the elasticity measured is the average elasticity that prevails in the population. The small difference can be due to a finite-sample bias affecting the estimation when applied to narrower samples.

5 Robustness checks

5.1 Depreciation of the permanent shock

I consider a more general income process in which permanent income is not necessarily a random walk, but simply an AR(1) with coefficient ρ:

$$p_t = \rho p_{t-1} + \eta_t$$

This still means that the permanent shock η_t affects the value of permanent income at each period in the future until the end of the household's life, but the effect of η_t now depreciates at a rate $(1-\rho)$ over time instead of affecting future permanent income in the same way at all future periods. Log-income growth now depends on all the past permanent shocks experienced by the household:

$$\Delta ln(y_t) = \eta_t - (1-\rho)p_{t-1} + \varepsilon_t - (1-\theta)\varepsilon_{t-1} - \theta\varepsilon_{t-2}$$

As a result, current log-income growth covaries with current log-consumption growth through the current transitory and permanent shocks, the past transitory shocks up to $t-2$, but also through p_{t-1}, which depends on all past permanent shocks. The instrument, future log-income growth at $t+2$, ca no longer identify separately the effect of the current transitory shock, because it now depends on η_t and p_{t-1}.

Kaplan and Violante (2010) show, however, that if the value of ρ is known, it is possible to obtain a consistent estimator in the presence of depreciation by substituting log-consumption growth for the

[18]The MPC of households in the low income group is 0.24, and it is 0.21 in the high income group, and only the second is significant at the 10% level. The MPC is not statistically different from zero among both households with low liquid wealth and high liquid wealth. The MPC of households with mortgage debt is 0.26 and the MPC of households who own a house without mortgage is 0.47, both statistically significant at the 5% level (Table 6 in their paper).

quasi-difference $\Delta^{\rho} ln(y_{i,t}) ln(y_t) - \rho ln(y_{t-1})$ in the estimating moments. The estimator then writes:

$$\hat{\phi}^{\varepsilon}_{\rho} = \frac{cov(\Delta ln(c_{i,t}), \Delta^{\rho} ln(y_{i,t+2}))}{cov(\Delta^{\rho} ln(y_{i,t}), \Delta^{\rho} ln(y_{i,t+2}))} = \phi^{\varepsilon}$$

Table 5: Estimated elasticities ϕ - Depreciation of the permanent shock at different rates $1-\rho$

Persistence	$\rho = 1$	$\rho = 0.96$	$\rho = 0.92$	$\rho = 0.88$
ϕ^{ε}	0.465***	0.487**	0.560**	0.775**
	(0.172)	(0.189)	(0.237)	(0.382)
Implied MPC	0.387***	0.406**	0.467**	0.646**
	(0.143)	(0.158)	(0.198)	(0.319)

Note: This table presents estimates of the average elasticity of nondurable consumption expenditures when measured with an estimator that incorporates a depreciation in the effect of permanent income shocks at a rate $1-\rho$.

Table 5 presents the results obtained with a depreciation-consistent estimator, for different values of ρ. The first column is the baseline case when $\rho = 1$. The point estimate of the elasticity of consumption then increases as the depreciation rate $1-\rho$ increases: it moves gradually from 0.465 to 0.487, 0.560, and 0.775 with decreases in ρ from 1 t 0.96, 0.92, and 0.88. It means that, if the true depreciation rate is smaller than one, my robust estimator is in fact conservative and underestimates the true elasticity.

5.2 Anticipation of the shocks

Anticipation of permanent shocks I examine how the estimator would be affected if part of the permanent income shock was anticipated by the household:

$$\eta_{i,t} = \eta^{s}_{i,t} + \eta^{a,t-s}_{i,t}$$

The permanent shock writes as the sum of a surprise component $\eta^{s}_{i,t}$ and an anticipated component $\eta^{a,t-s}_{i,t}$ that is in the information set of the consumer at $t-s$ but only realizes at t. The anticipated component is uncorrelated with all other surprise shocks and anticipated shocks. In such a model, the robust estimator I use takes the following expression:

$$\begin{aligned}\hat{\phi}^{\varepsilon} &= \frac{cov(\Delta ln(c_{i,t}), \Delta^{\rho} ln(y_{i,t+2}))}{cov(\Delta^{\rho} ln(y_{i,t}), \Delta^{\rho} ln(y_{i,t+2}))} \\ &= \frac{-\theta cov(\Delta ln(c_{i,t}), \varepsilon_{i,t}) + \overbrace{cov(\Delta ln(c_{i,t}), \eta^{a,t+2-s}_{i,t+2})}^{=0 \text{ if } s=1 / >0 \text{ if } s=2}}{-\theta var(\varepsilon_{i,t})} = \phi^{\varepsilon} \text{ if } s=1 \ / \ < \phi^{\varepsilon} \text{ if } s=2\end{aligned}$$

Because the shocks are uncorrelated, the denominator $cov(\Delta^{\rho} ln(y_{i,t}), \Delta^{\rho} ln(y_{i,t+2}))$ is unaffected by the presence of an anticipated component, regardless of when it is anticipated. However, receiving income news or having received income news can affect log-consumption growth depending on how early the shock is anticipated. When the anticipated component is observed one period in advance, then the permanent shock at $t+2$ is partly anticipated at $t+1$ but remains uncorrelated with log-consumption growth at t. The robust estimator remains consistent. When the anticipated component of a permanent shock is observed two periods in advance, then the permanent shock at $t+2$ is partly anticipated at t, and the realization of the anticipated component correlates positively with log-consumption growth at t. This means that the robust estimator is conservative, and underestimates the true elasticity: the anticipation induces a positive correlation between current consumption growth and future log-income growth while the response of consumption to a current transitory shock generates a negative correlation between the two. Not controlling for the anticipation induces a downward bias.

Anticipation of transitory shocks I assume that the transitory shock decomposes into a surprise component and an anticipated component:

$$\varepsilon_{i,t} = \varepsilon^{s}_{i,t} + \varepsilon^{a,t-s}_{i,t}$$

The estimator becomes:

$$\begin{aligned}\hat{\phi}^{\varepsilon} &= \frac{cov(\Delta ln(c_{i,t}), \Delta^{\rho} ln(y_{i,t+2}))}{cov(\Delta^{\rho} ln(y_{i,t}), \Delta^{\rho} ln(y_{i,t+2}))} \\ &= \frac{-\theta cov(\Delta ln(c_{i,t}), \varepsilon_{i,t}) + \overbrace{cov(\Delta ln(c_{i,t}), \varepsilon^{a,t+2-s}_{i,t+2} - (1-\theta)\varepsilon^{a,t+1-s}_{i,t+1})}^{<0 \text{ if } s=1}}{-\theta var(\varepsilon_{i,t})} > \phi^{\varepsilon} \text{ if } s=1\end{aligned}$$

The denominator of the ratio is unaffected by the partial anticipation, but the nominator incorporates new terms. In particular, when the transitory income shock is anticipated one period in advance, then part of the negative correlation observed between log-consumption growth and future log-income growth and attributed to the response of consumption to the current shock could be partly driven by the response of consumption to the anticipated component of the future transitory shock at $t+1$. Note that, this effect can only be large if consumption responds strongly to news about transitory income, that is , if it responds strongly to transitory shocks. When the transitory income shock is anticipated two periods in advance, and the effect of current news on log-consumption growth is stronger than the effect of past news, $cov(\Delta ln(c_{i,t}), \varepsilon^{a,t}_{i,t+2}) > cov(\Delta ln(c_{i,t}), (1-\theta)\varepsilon^{a,t-1}_{i,t+1})$, the direction of the bias is reversed. Thus, the ground for the robust estimate overstating the true elasticity because of partial anticipation seems limited.

5.3 Clustering

Table 6: Estimated elasticities ϕ - Different clustering

Clustering	year $\times$ educ	household	year
ϕ^{ε}	0.465***	0.465*	0.465**
	(0.172)	(0.282)	(0.219)
Implied MPC	0.387***	0.387*	0.387**
	(0.143)	(0.235)	(0.183)
Obs.	8,958	8,958	8,958

Note: This table presents estimates of the average elasticity of non-durable consumption expenditures for different choices of clustered estimation residuals.

Although the estimator is unbiased, it is likely that there is some correlation in the estimation residuals in the sample. This would alter the measured precision of the estimation. In the baseline specification, I assume that at a given year households with a similar occupation can be subject to similar shocks, and I cluster by a year $\times$ education level dummy to control for this effect. This corresponds to 27 clusters. The baseline results are presented in the first column as a reference point. I consider here the possibility of other types of correlations. The second column examine the precision of the estimation when clustering at the household level, that is, allowing for a correlation between the estimation residual of a household over time. That represents 1561 clusters. The precision decreases, but the point estimate remains statistically significant at the 10% level. The third and fourth columns show the results associated with a clustering only at the year or at the education level, that is, allowing for a correlation between the estimation residuals of all households at a given year (9 clusters). The point estimate remains statistically significant at the 5% level.

5.4 Change in the set of detrending variables

Table 7: Estimated elasticities ϕ - Different sets of detrending variables

Detrending	Baseline	+ extra (BPP)	- empl. stat.	- res. dum.	- fam dum.	year/cohort
ϕ^{ε}	0.465***	0.424**	0.488*	0.344*	0.418**	0.437
	(0.172)	(0.185)	(0.254)	(0.196)	(0.203)	(0.297)
Implied MPC	0.388***	0.354**	0.407*	0.287*	0.348**	0.365
	(0.143)	(0.154)	(0.212)	(0.164)	(0.169)	(0.248)
Obs.	8,958	8,958	8,958	8,958	8,958	

Note: This table presents estimates of the average elasticity of nondurable consumption expenditures for different sets of detrending variables. The first column corresponds to the baseline set. The second column adds to the baseline set a dummy for having an extra income recipient besides the male and female head in the households. The third column removes from the baseline set the dummy for employment status (emplyed, unemployed, retired). The fourth column removes from the baseline set the residential dummies, that is, a region dummy and a dummy for whether the household lives in a big city. The fourth columns removes from the baseline set the family dummies, that is, the number of children, number of household members, and number of children out of the household. The fifth column uses only the year and year of birth of the male head as detrending variables.

Table 8 presents the results obtained when varying the set of variables used to detrend log-consumption and log-income growth. It shows that the elasticity remains large across all sets of variables, and no point estimate is statistically different from another at the 10% level. This is consistent with a model in which demographic characteristics are unaffected by contemporaneous transitory shocks, so that detrending or not does not modify the point estimates. The estimates are all statistically different from at the 10% level, except when detrending only with the year and year of birth of the male head. The first column corresponds to the baseline results. The second column shows that the results remain very similar when adding a dummy for the presence of an extra income recipient in the household, as do BPP—the reason I do not is to get more variability by not controlling for income shocks on the revenue of other members of the household. The point estimate is a little smaller, suggesting that experiencing a good transitory shock reduces a little the likelyhood of having an extra income recipient in the household and the extra consumption expenditures that correlate with it. Similarly, the point estimate increases a little when removing the employment status variables, suggesting that a positive transitory shock correlates with a change in status that is positively associated with consumption expenditures. The point estimate decreases when removing the residential dummies, suggesting that a positive transitory shock correlates with a change in residence that is negatively associated with consumption expenditures. The point estimate decreases also a little when removing the family dummies, suggesting that a positive transitory shock correlates with a change in family composition that is negatively associated with consumption expenditures. Finally, the point estimate remains very similar though not precisely estimated when using only the year and year of birth dummies as detrending variables. It is important to keep those to control for correlation through time trends in the moments used for estimation.

6 Conclusion

This paper shows that the discrepancy between the response of consumption to a transitory shock estimated in natural experiments and the values derived from structural estimation methods is resolved when allowing for a correlation between log-consumption growth and past shocks in structural approaches.

I show that such a correlation exists, because of precautionary behavior, in the life-cycle model underlying structural methods. Indeed, precautionary behavior simultaneously responds to the value of past shocks and contributes to the value of log-consumption growth. Also, other features commonly incorporated in life-cycle models, including exogenous borrowing constraints, wealthy-hand-to-mouth behavior and habit persistence, generate a similar correlation.

I prove that this correlation induces a downward bias in structural estimators. This is because, in these methods, the negative impact of past shocks on log-consumption growth partly offsets the positive impact of current transitory shocks, causing an underestimation of the consumption elasticity to a transitory shock. On the contrary, the negative impact of past shocks on log-consumption growth increases the value of the restriction used to identify the consumption elasticity to a permanent shock, leading to an overestimation of this elasticity.

When adapting the estimating restrictions to take into account the impact of past shocks on log-consumption growth, the elasticity to a transitory shock becomes statistically significant and the point estimate jumps from 0.07 to 0.47. The MPC associated with such an elasticity is 39%, meaning that 39% of a transitory income gain would be consumed over the following year. This figure matches with the evidence obtained from natural experiments, which place the quarterly MPC between 9% and 38%—with the response of consumption over later quarters not being precisely estimated so that it could compare with a yearly MPC. The elasticity of consumption to a past transitory shock is negative, as predicted by the theory, and statistically significant at the 10% level. Considering separately households with different levels of liquidity, income, education level and age, the elasticity of consumption to a transitory shock remains statistically significant across most types of households, even those with high financial income and a high wage rate that are not likely to be constrained. This matches with studies based on natural experiments that look at the heterogeneity of the response across households.

These results give external validity to responses measured from idiosyncratic set-ups: households do not respond to a tax rebate only because they might expect it to be permanent, they do not respond to a fiscal stimulus simply because those stimulus often occur in times of recession, they do not respond to a lottery gain only in Norway or to a pay freeze related to the government shut down only because they are Federal government employees usually subject to little income uncertainty that have not accumulated liquidity. It also provides a simple and robust method to be applied to new data in the future.

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 Detailed derivations on precautionary behavior

B.1 Log-consumption growth

The Euler equation is:

$$\begin{aligned} u'(c_t)e^{\delta_t z_t} &= E_t[u'(c_{t+1})]\beta(1+r)e^{\delta_{t+1}z_{t+1}} \\ u'(c_t) &= E_t[u'(c_{t+1})]\beta(1+r)e^{\delta_{t+1}z_{t+1}-\delta_t z_t} \\ u'(c_t) &= E_t[u'(c_{t+1})]R_{t,t+1}. \end{aligned} \tag{B.1}$$

with $R_{t,t+1} = \beta(1+r)e^{\delta_{t+1}z_{t+1}-\delta_t z_t}$. Following Kimball (1990), I define φ_t the equivalent precautionary premium associated with risk on consumption at $t+1$ from the point of view of period t. It is the variable φ_t such that:

$$E_t\left[u'(c_{t+1})\right] = u'(E_t[c_{t+1}] - \varphi_t). \tag{B.2}$$

From Jensen's inequality, the premium φ_t is strictly positive for strictly prudent consumers (marginal utility is strictly convex) and zero for certainty-equivalent consumers (marginal utility is linear)[19]

Combining this expression for future expected marginal utility (A.2) with the Euler equation (A.1) yields:

$$\begin{aligned} u'(c_t) &= u'(E_t[c_{t+1}] - \varphi_t)R_{t,t+1} \\ c_t^{-\rho} &= (E_t[c_{t+1}] - \varphi_t)^{-\rho}R_{t,t+1} \\ c_t &= (E_t[c_{t+1}] - \varphi_t)R_{t,t+1}^{-1/\rho} \\ E_t[c_{t+1}] &= c_t R_{t,t+1}^{1/\rho} + \varphi_t \end{aligned} \tag{A.3}$$

I express $E_t[c_{t+1}] = c_{t+1} - \xi_{t+1}$, and divide each side by $c_t R_{t,t+1}^{1/\rho}$:

$$\begin{aligned} c_{t+1} - \xi_{t+1} &= c_t R_{t,t+1}^{1/\rho} - \varphi_t \\ \frac{c_{t+1}}{c_t R_{t,t+1}^{1/\rho}} &= 1 + \frac{\varphi_t + \xi_{t+1}}{c_t R_{t,t+1}^{1/\rho}} \end{aligned}$$

I take the logarithm of each side:

$$\begin{aligned} \Delta ln(c_{t+1}) &= ln(R_{t,t+1}^{1/\rho}) + ln\left(1 + \frac{\varphi_t + \xi_{t+1}}{c_t R_{t,t+1}^{1/\rho}}\right) \\ \Delta ln(c_{t+1}) &= \frac{1}{\rho} ln(\beta(1+r)) + \frac{1}{\rho}\Delta\delta_{t+1} z_{t+1} + ln\left(1 + \frac{\varphi_t + \xi_{t+1}}{c_t R_{t,t+1}^{1/\rho}}\right) \end{aligned}$$

I take a Taylor approximation around $(\varepsilon_{t+1}, \eta_{t+1}) = (0,0)$, denoting with an exponent 0 the variables at this point. The only variable to depend on the value of ε_{t+1} and η_{t+1} is the consumption innovation, ξ_{t+1}. Therefore it writes:

$$\begin{aligned} \Delta ln(c_{t+1}) = {} & \frac{1}{\rho} ln(\beta(1+r)) + \frac{1}{\rho}\Delta\delta_{t+1} z_{t+1} + ln\left(1 + \frac{\varphi_t + \xi_{t+1}^0}{c_t R_{t,t+1}^{1/\rho}}\right) \\ & + \varepsilon_{t+1}\frac{d\Delta ln(c_{t+1})}{d\varepsilon_{t+1}} + \eta_{t+1}\frac{d\Delta ln(c_{t+1})}{d\eta_{t+1}} + o(\varepsilon_{t+1}, \eta_{t+1}) \end{aligned} \tag{A.4}$$

The term $\xi_{t+1}^0 = c_{t+1}^0 - E_t[c_{t+1}]$, which is innovation to consumption at $(\varepsilon_{t+1}, \eta_{t+1}) = (0,0)$, is strictly positive because the partial derivatives of consumption with respect to ε and η are strictly negative, making the consumption functions of each shock concave (Commault (2016)). As a result, denoting

[19] Indeed, using Jensen's inequality, the convexity of $u'(c)$ implies:

$$E_t\left[u'(c_{t+1})\right] > u'(E_t[c_{t+1}]) \;\Leftrightarrow\; u'(E_t[c_{t+1}] - \varphi_t) > u'(E_t[c_{t+1}]) \;\Leftrightarrow\; E_t[c_{t+1}] - \varphi_t < E_t[c_{t+1}] \;\Leftrightarrow\; 0 < \varphi_t.$$

consumption as a function of the shocks as follows, $c_{t+1}(\varepsilon_{t+1},\eta_{t+1})$, I obtain that:

$$\begin{aligned}
&c_{t+1}(0,0) = c_{t+1}(E_t[\varepsilon_{t+1}],E_t[\eta_{t+1}]) \\
&c_{t+1}(0,0) > E_t[c_{t+1}(\varepsilon_{t+1},E_t[\eta_{t+1}])] \\
&c_{t+1}(0,0) > E_t[E_t[c_{t+1}(\varepsilon_{t+1},\eta_{t+1})] \\
&c_{t+1}(0,0) > E_t[c_{t+1}]
\end{aligned}$$

One moves from the first to the second line by seeing that because, when η is set at $E_t[\eta_{t+1}]$, the consumption function of ε is concave in ε. Ones goes from the second to the third by seeing that, fixing ε_{t+1}, consumption is concave in η and $c_{t+1}(\varepsilon_{t+1},E_t[\eta_{t+1}]) > E_t[c_{t+1}(\varepsilon_{t+1},\eta_{t+1})$. Because this is true for at any ε_{t+1}, it is true on average. Under perfect foresight, this term is zero because c_{t+1} is an exogenous fraction of total expected resources at $t+1$, $\left((1+r)a_t + p_t e^{\eta_{t+1}} e^{\varepsilon_{t+1}} \sum_{s=0}^{T-t} \frac{p_t e^{\eta_{t+1}}}{(1+r)^s}\right)$, which coincides with its expected value at t when the income shocks are equal to their expected value $(\varepsilon_{t+1},\eta_{t+1}) = (0,0)$. Therefore, the presence of the term $\xi^0_{t+1} > 0$ results from precautionary behavior too

B.2 Theorem

At any period $0 < t < T$, and for any $0 < k < t$:

$$\frac{d\varphi_t}{da_t} < 0 \text{ and } \frac{d\varphi_t}{d\varepsilon_{t-k}} < 0.$$

Response of φ_t to a change in net assets First, I derive each side of the equilibrium relation (A.3) with respect to a_t:

$$\frac{dE_t[c_{t+1}]}{da_t} = \frac{dc_t}{da_t} R^{1/\rho}_{t,t+1} + \frac{d\varphi_t}{da_t} \tag{A.5}$$

Because φ_t is the difference between expected future consumption and current consumption, when expected consumption $E_t[c_{t+1}]$ reacts less than current consumption c_t, then the precautionary premium has decreased, and vice-versa.

To compare the response of consumption over time, I derive each side of the Euler equation with respect to a_t:

$$\begin{aligned}
&\frac{dc_t}{da_t} u''(c_t) = E_t[\frac{dc_{t+1}}{da_t} u''(c_{t+1})] R_{t,t+1} \\
&\frac{dc_t}{da_t} R^{1/\rho}_{t,t+1} = E_t[\frac{dc_{t+1}}{da_t} \frac{u''(c_{t+1})}{u''(c_t)}] R^{(1+\rho)/\rho}_{t,t+1} \\
&\frac{dc_t}{da_t} R^{1/\rho}_{t,t+1} = E_t[\frac{dc_{t+1}}{da_t}] \underbrace{E_t[\frac{u''(c_{t+1})}{u''(c_t) R^{-(1+\rho)/\rho}_{t,t+1}}]}_{>1} + \underbrace{cov_t\left(\frac{dc_{t+1}}{da_t}, \frac{u''(c_{t+1})}{u''(c_t)}\right) R^{(1+\rho)/\rho}_{t,t+1}}_{>0}
\end{aligned} \tag{A.6}$$

To establish that $E_t[\frac{u''(c_{t+1})}{u''(c_t) R^{-(1+\rho)/\rho}_{t,t+1}}] > 1$, I use the fact that, when $\frac{-u'''(c)}{u''(c)}$ is decreasing, which is the case under isoelastic preferences, the second derivative of marginal utility is a convex function of marginal utility[20]. I denote $-u''(c) = f(u'(c))$, with f strictly convex. It implies:

$$\begin{aligned}
E_t[-u''(c_{t+1})] &= E_t[f(u'(c_{t+1}))] \\
E_t[-u''(c_{t+1})] &> f\left(E_t[u'(c_{t+1})]\right) = f\left(u'(c_t)R_{t,t+1}^{-1}\right) = f\left(u'(c_t R_{t,t+1}^{1/\rho})\right) \\
E_t[-u''(c_{t+1})] &> -u''(c_t R_{t,t+1}^{1/\rho}) = -u''(c_t)R_{t,t+1}^{-1+\rho)/\rho} \\
\frac{E_t[-u''(c_{t+1})]}{-u''(c_t)R_{t,t+1}^{-1+\rho)/\rho}} &> 1
\end{aligned}$$

I pass from the first to the second line using the fact that $f(c)$ is strictly convex, so that $E[f(c)] > f(E[c])$. Within the second line, I use the Euler equation to have that $E_t[u'(c_{t+1})] = u'(c_t)R_{t,t+1}^{-1}$. Finally,to move on from the second to the third line, I use that $f(u'(c)) = -u''(c)$.

To establish that the covariance term is strictly positive, I rely on a result of Commault (2016), showing that $\frac{d^2c_{t+1}}{da_{t+1}\varepsilon_{t+1}} < 0$ and $\frac{d^2c_{t+1}}{da_{t+1}\eta_{t+1}} < 0$ (and also the more simple result that $\frac{da_{t+1}}{da_t} > 0$). Therefore:

$$\begin{cases}
\frac{d^2c_{t+1}}{da_t\varepsilon_{t+1}} = \underbrace{\frac{da_{t+1}}{da_t}}_{>0}\underbrace{\frac{d^2c_{t+1}}{da_{t+1}\varepsilon_{t+1}}}_{<0} < 0 \\
\frac{d^2c_{t+1}}{da_t\eta_{t+1}} = \underbrace{\frac{da_{t+1}}{da_t}}_{>0}\underbrace{\frac{d^2c_{t+1}}{da_{t+1}\eta_{t+1}}}_{<0} < 0.
\end{cases}$$

Regarding the variations in $-u''(c_{t+1})$, Commault (2016) proves the simple results that $\frac{dc_{t+1}}{d\varepsilon_{t+1}} > 0$ and $\frac{dc_{t+1}}{d\eta_{t+1}} > 0$. Therefore:

$$\begin{cases}
\frac{d(-u''(c_{t+1}))}{d\varepsilon_{t+1}} = -\underbrace{u'''(c_{t+1})}_{>0}\underbrace{\frac{dc_{t+1}}{d\varepsilon_{t+1}}}_{>0} < 0 \\
\frac{d(-u''(c_{t+1}))}{d\eta_{t+1}} = -\underbrace{u'''(c_{t+1})}_{>0}\underbrace{\frac{dc_{t+1}}{d\eta_{t+1}}}_{>0} < 0.
\end{cases}$$

Because the only two shocks that occur between t and $t+1$ move each side of the covariance in the same

[20] Intuitively, this condition means that $-u'''(c)$, the change in $-u''(c)$, increases faster than $u''(c)$, the change in $u'(c)$, thus that $-u''(c)$ is more convex than $u'(c)$. Formally, this conditions implies that the coefficient of risk aversion, which measures the convexity of a function, is always larger for $-u''(c)$ than for $u'(c)$:

$$\begin{aligned}
\frac{d\frac{u'''(c)}{-u''(c)}}{dc} < 0 &\Leftrightarrow \frac{u'''(c)}{-u''(c)}\left(\frac{u'''(c)}{-u''(c)} - \frac{-u''''(c)}{u'''(c)}\right) < 0 \\
&\Leftrightarrow \frac{u'''(c)}{-u''(c)} < \frac{-u''''(c)}{u'''(c)} \\
&\Leftrightarrow \pi_{u'} < \pi_{-u''} \\
&\Leftrightarrow -u''(c) \text{ is a convex function of } u'(c)
\end{aligned}$$

The last line is obtained from Arrow (1965) and Pratt (1964), who show that if the risk premium of a function is strictly larger than that of an another function for any value of c, then the former is a convex transformation of the latter.

direction, the covariance is positive:

$$cov_t\left(\frac{dc_{t+1}}{da_t}, -u''(c_{t+1})\right) > 0.$$

Then, combining equations (A.5) and (A.6) yields:

$$\frac{d\varphi_t}{da_t} = \frac{dE_t[c_{t+1}]}{da_t} - \frac{dc_t}{da_t}R_{t,t+1}^{1/\rho} < 0.$$

Response of φ_t to a change in a past transitory shock Because past transitory shocks only affect the precautionary through their impact on net assets, this means that past transitory shocks are negatively correlated with the premium φ_t. For any $k > 1$:

$$\frac{d\varphi_t}{d\varepsilon_{t-k}} = \underbrace{\frac{da_t}{d\varepsilon_{t-k}}}_{>0}\underbrace{\frac{d\varphi_t}{da_t}}_{<0} < 0$$

Indeed, the premium is entirely determined from the expected distribution of future consumption c_{t+1}. Now, consumption is chosen as the solution of problem (2.1)-(2.3), which means it only depends on the household's level of net assets, permanent income, immediately past transitory shock (because this shock lasts for one period thus influences the level of their current income) and current transitory shock. The only item in this list that transitory shocks at $t-2$ and before affect is net assets.

To determine the impact of past transitory shocks on current assets, I use the Euler equation: it implies that the response of expected marginal utility to a transitory shock must be the same at all periods. As a result, consumption must respond in the same direction at all periods, and assets too, so that the change in transitory income to be transmitted to later periods. Formally, I derive each side of the Euler equation with respect to ε_{t-1}:

$$\underbrace{\frac{dc_{t-1}}{d\varepsilon_{t-1}}}_{>0}\underbrace{(-u''(c_{t-1}))}_{>0} = \frac{da_t}{d\varepsilon_{t-1}}E_t[\underbrace{\frac{dc_t}{da_t}}_{>0}\underbrace{(-u''(c_t))}_{>0}]$$

Because all terms are strictly positive, $\frac{da_t}{d\varepsilon_{t-1}}$ has to be strictly positive too. Similarly, a change in past net assets has a positive impact on current net assets. Iterating backwards, a past transitory shock ε_{t-k} raises a_{t-k+1} and then raises assets at each later period, including a_t:

$$\frac{da_t}{d\varepsilon_{t-k}} > 0$$

Regarding transitory shocks at $t-1$, they affect φ_t both their impact on a_t and on y_t. To prove that they are negatively correlated with past shocks, I apply the same proof as used for net assets:

$$\frac{dc_t}{d\varepsilon_{t-1}}u''(c_t) = E_t[\frac{dc_{t+1}}{d\varepsilon_{t-1}}u''(c_{t+1})]R_{t,t+1}$$
$$\frac{dc_t}{d\varepsilon_{t-1}}R^{1/\rho}_{t,t+1} = E_t[\frac{dc_{t+1}}{d\varepsilon_{t-1}}]\underbrace{E_t[\frac{u''(c_{t+1})}{u''(c_t)R^{-(1+\rho)/\rho}_{t,t+1}}]}_{>1} + \underbrace{\frac{da_{t+1}}{d\varepsilon_{t-1}}cov_t\left(\frac{dc_{t+1}}{da_{t+1}}, \frac{u''(c_{t+1})}{u''(c_t)}\right)R^{(1+\rho)/\rho}_{t,t+1}}_{>0}$$

The term $E_t[\frac{u''(c_{t+1})}{u''(c_t)R^{-(1+\rho)/\rho}_{t,t+1}}]$ is larger than one for the same reason as above. The covariance is negative too because, as shown above, $\frac{da_{t+1}}{d\varepsilon_{t-1}} > 0$. Thus:

$$\frac{d\varphi_t}{d\varepsilon_{t-1}} = \frac{dE_t[c_{t+1}]}{d\varepsilon_{t-1}} - \frac{dc_t}{d\varepsilon_{t-1}}R^{1/\rho}_{t,t+1} < 0$$

B.3 Corollary

In the model presented above, log-consumption growth is strictly decreasing in the value of past transitory shocks:

$$cov_t(\Delta ln(c_{t+1}), \varepsilon_{t-k}) < 0.$$

Substituting for $\Delta ln(c_{t+1})$ with the expression (A.4) I break down the log-consumption growth into its different elements:

$$cov_t(\Delta ln(c_{t+1}), \varepsilon_{t-k}) = \underbrace{cov_t(\Delta ln(R^{1/\rho}_{t,t+1}), \varepsilon_{t-k})}_{=0} + cov_t(ln(1 + \frac{\varphi_t + \xi^0_{t+1}}{c_t R^{1/\rho}_{t,t+1}}), \varepsilon_{t-k})$$
$$+ \underbrace{E_t[\varepsilon_{t+1}]}_{=0} cov_t(\frac{d\Delta ln(c_{t+1})}{d\varepsilon_{t+1}}, \varepsilon_{t-k}) + \underbrace{E_t[\eta_{t+1}]}_{=0} cov_t(\frac{d\Delta ln(c_{t+1})}{d\eta_{t+1}}, \varepsilon_{t-k})$$

The covariance is therefore driven by the precautionary component of log-consumption growth, $ln(1 + \frac{\varphi_t + \xi^0_{t+1}}{c_t R^{1/\rho}_{t,t+1}})$. I have established above that $\frac{d\varphi_t}{d\varepsilon_{t-k}} < 0$. Also, within the course of the proof, I have shown that $\frac{dc_t}{d\varepsilon_{t-k}} = \frac{dc_t}{da_t}\frac{da_t}{d\varepsilon_{t-k}} > 0$. Let me know study the response of the consumption innovation ξ_{t+1} at the point $(0,0)$. I show in Commault (2016) that the partial derivative of $\frac{dc_{t+1}}{da_t}$ with respect to ε and to η are strictly positive, so that the functions $\frac{dc_{t+1}}{da_t}(\varepsilon)$ and $\frac{dc_{t+1}}{da_t}(\eta)$ are convex. As a result:

$$\begin{aligned}\left(\frac{d\xi_{t+1}}{da_t}\right)^0 &= \frac{dc^0_{t+1}}{da_t} - \frac{dE_t[c_{t+1}]}{da_t} \\ &< E_t[\frac{dc^0_{t+1}}{da_t}] - \frac{dE_t[c_{t+1}]}{da_t} \\ &< 0\end{aligned}$$

Finally:

$$\left(\frac{d\xi_{t+1}}{d\varepsilon_{t-k}}\right)^0 = \underbrace{\frac{da_t}{d\varepsilon_{t-k}}}_{>0}\underbrace{\left(\frac{d\xi_{t+1}}{da_t}\right)^0}_{<0} < 0$$

As a result the relation between the precautionary component of log-consumption growth and past shocks is strictly negative:

$$\frac{dln(1+\frac{\varphi_t+\xi^0_{t+1}}{c_t R^{1/\rho}_{t,t+1}})}{d\varepsilon_{t-k}} = \left(\underbrace{\left(\frac{d\varphi_t}{d\varepsilon_{t-k}} + \left(\frac{d\xi_{t+1}}{d\varepsilon_{t-k}}\right)^0\right)}_{<0}\frac{1}{c_t R^{1/\rho}_{t,t+1}} - \underbrace{\frac{dc_t R^{1/\rho}_{t,t+1}}{d\varepsilon_{t-k}}}_{>0}\frac{\varphi_t+\xi^0_{t+1}}{c_t R^{1/\rho}_{t,t+1}}\right)\frac{1}{c_t R^{1/\rho}_{t,t+1}+\varphi_t+\xi^0_{t+1}} < 0$$

B.4 Consumption elasticities

I look into the implications of precautionary behavior for the consumption elasticities ϕ^ε_t and ϕ^η_t. To do so, I first express consumption and precautionary saving in levels rather than in growth terms. Iterating forward on equation (2.6), consumption growth between t and any future period $t+s$ is a weighted sum of the precautionary premiums between these two dates:

$$E_t[c_{t+s}] = c_t R^{1/\rho}_{t,t+s} + \underbrace{\sum_{k=1}^{s} E_t[\varphi_{t+k-1}] R^{1/\rho}_{t+k,t+s}}_{\substack{\text{precautionary}\\ \text{consumption growth}\\ \text{between } t \text{ and } t+s}}.$$

What is the additional amount of saving at t necessary to implement this precautionary consumption growth between t and all future periods? [21] I plug the equation above into the intertemporal budget constraint (2.4), to obtain the following equilibrium relationship satisfied by consumption:

$$c_t = \underbrace{\frac{1}{l_{t,0}}\left((1+r)a_t + \sum_{s=0}^{T-t}\frac{E_t[y_{t+s}]}{(1+r)^s}\right)}_{\substack{\text{consumption under perfect foresight}\\ \frac{1}{l_{t,0}}W_t}} - \underbrace{\frac{1}{l_{t,0}}\left(\sum_{s=1}^{T-t}\frac{l_{t,s}E_t[\varphi_{t+s-1}]}{(1+r)^s}\right)}_{\substack{\text{precautionary saving}\\ \frac{1}{l_{t,0}}PG_t}}. \tag{2.7}$$

Under perfect foresight, the household consumes a fraction $\frac{1}{l_{t,0}}$ of its total expected lifetime resources—the sum of its net assets, current income and expected future income—denoted W_t (for wealth). By definition, precautionary saving is the difference between actual consumption and the level of consumption that would be chosen under perfect foresight. This equilibrium relationship shows that this difference corresponds to a fraction $\frac{1}{l_{t,0}}$ of the household's expected lifetime precautionary consumption growth, denoted PG_t (for precautionary growth): to be able to implement the consumption growth it desires, the household takes out its total expected precautionary growth from its total expected resources and consumes a constant share of what remains. Thus, the behavior of a household facing risk can be inter-

[21]Note that simply saving an additional amount φ_k at each period k is not the solution: this would reduce consumption at t by φ_t, increase consumption at $t+1$ by a fraction of this additional saving, and reduce consumption at $t+1$ by φ_{t+1}. The resulting consumption growth has no reason to coincide with φ_k in general.

preted as a permanent-income style decision, but applied to an uncertainty-adjusted measure of its total expected resources, instead of its total expected resources (as under perfect foresight). The exogenous fraction $\frac{1}{l_{t,0}} = (\sum_{s=0}^{T-t} \frac{R_{t,t+s}^{1/\rho}}{(1+r)^s})^{-1}$ measures the weight put on consumption at period t within their total lifetime consumption[22]. More generally, the fraction $\frac{1}{l_{t,s}} = (\sum_{k=0}^{T-t-s} \frac{R_{t+s,t+s+k}^{1/\rho}}{(1+r)^k})^{-1}$ represents the weight put on consumption between the beginning of period t and the beginning of period $t+s+1$.

Using equation (2.7), it is then possible to explicit the contribution of precautionary behavior for the consumption elasticities to the income shocks ε and η:

$$\phi_t^{\varepsilon} = \frac{\left(\frac{dW_{t+1}}{d\varepsilon_{t+1}}\right)^0 - \overbrace{\left(\frac{dPG_{t+1}}{d\varepsilon_{t+1}}\right)^0}^{\text{precaution (1)}}}{W_{t+1} - \underbrace{(PG_{t+1})^0}_{\text{precaution (2)}}} \quad \text{and} \quad \phi_t^{\eta} = \frac{\left(\frac{dW_{t+1}}{d\eta_{t+1}}\right)^0 - \overbrace{\left(\frac{dPG_{t+1}}{d\eta_{t+1}}\right)^0}^{\text{precaution (1)}}}{W_{t+1} - \underbrace{(PG_{t+1})^0}_{\text{precaution (2)}}}.$$

with $W_{t+1} = (1+r)a_{t+1} + \sum_{s=0}^{T-t-1} \frac{E_{t+1}[y_{t+1+s}]}{(1+r)^s}$ the total expected resources at $t+1$, and $PG_{t+1} = \sum_{s=1}^{T-t-1} l_{t+1,s} \frac{E_{t+1}[\varphi_{t+s}+\tilde{\lambda}_{t+s}]}{(1+r)^s}$ the sum of future expected precautionary growth at date $t+1$. Precautionary behavior affects the elasticity of consumption in two ways, indicated with (1) and (2). First, denoted with (1), income shocks affect the amount of precautionary saving that the household needs to make, which can attenuate of reinforce their consumption response: if the shock increases income but raises precautionary saving, consumption responds less than it would under perfect foresight, while if the shock increases income and reduces precautionary saving, consumption responds more. Second, denoted with (2), because the household saves for precautionary reasons, it consumes less at $t+1$ than they would under perfect foresight, so that a given change in consumption corresponds to a larger percentage change in consumption for consumers engaging in precautionary behavior than for consumers with perfect foresight: the denominator of the elasticity of consumption is always smaller.

B.5 Precautionary behavior in BPP

I follow the derivation presented in Blundell, Low, and Preston (2013), which is identical but more detailed than in BPP. The authors begin with the Euler equation:

$$E_t[u'(c_{t+1})] = u'(c_t e^{R_{t,t+1}^{1/\rho}})$$

Using an exact Taylor expansion around $e_{t+1} = 0$, the authors obtain that there exists a $\tilde{c}$, between $c_t R_{t,t+1}^{1/\rho}$ and c_{t+1} such that:

$$E_t[\Delta ln(c_{t+1})] = ln(R_{t,t+1}^{1/\rho}) + \frac{1}{2}\frac{u'(c_t R_{t,t+1}^{1/\rho})}{c_t R_{t,t+1}^{1/\rho} u''(c_t R_{t,t+1}^{1/\rho})} E_t[\frac{\tilde{c}u'''(\tilde{c}) + \tilde{c}u''(\tilde{c})}{u'(c_t R_{t,t+1}^{1/\rho})} e_{t+1}^2]$$

[22] When consumers are neither patient nor impatient ($\beta = \frac{1}{1+r}$) and individual characteristics are constant ($z_t = z$), $l_{t,0}$ tends toward $\frac{r}{1+r}$ as T approaches infinity.

with my notations (this equation corresponds to the unnumbered at the top of page 31 in Blundell, Low, and Preston (2013)). They decompose the expected value of $E_t[\Delta ln(c_{t+1})]$ as the sum of $\Delta ln(c_{t+1})$ plus an innovation, which is the term u_{t+1} with my notations:

$$\Delta ln(c_{t+1}) = ln(R_{t,t+1}^{1/\rho}) + \frac{1}{2}\underbrace{\frac{u'(c_t R_{t,t+1}^{1/\rho})}{c_t R_{t,t+1}^{1/\rho} u''(c_t R_{t,t+1}^{1/\rho})}}_{\gamma(c_t R_{t,t+1}^{1/\rho})} E_t[\underbrace{\frac{\tilde{c}u'''(\tilde{c}) + \tilde{c}u''(\tilde{c})}{u'(c_t R_{t,t+1}^{1/\rho})}}_{\beta(\tilde{c}, c_t R_{t,t+1}^{1/\rho})} e_{t+1}^2] + u_{t+1}$$

Now, because they erroneously assimilate e_{t+1} and u_{t+1}, they claim that the term $\gamma(c_t R_{t,t+1}^{1/\rho}) \times E_t[\beta(\tilde{c}, c_t R_{t,t+1}^{1/\rho}) e_{t+1}^2]$ is a $\mathcal{O}(E_t[u_{t+1}^2])$ while it is a $\mathcal{O}(E_t[e_{t+1}^2])$. They obtain:

$$\Delta ln(c_{t+1}^{BPP}) = ln(R_{t,t+1}^{1/\rho}) + \mathcal{O}(E_t[u_{t+1}^2]) + u_{t+1},$$

When no such assimilation is made it writes:

$$\begin{aligned} \Delta ln(c_{t+1}) &= ln(R_{t,t+1}^{1/\rho}) + \mathcal{O}(E_t[e_{t+1}^2]) + u_{t+1} \\ \Delta ln(c_{t+1}) &= ln(R_{t,t+1}^{1/\rho}) + \underbrace{\mathcal{O}(E_t[(u_{t+1} + E_t[ln(1 + \frac{\varphi_t + \xi_{t+1}}{R_{t,t+1}^{1/\rho}})])^2]}_{\neq \mathcal{O}(E_t[u_{t+1}^2])} + u_{t+1} \end{aligned}$$

The trend of log-consumption growth is *not deterministic* in this case, because the term $\mathcal{O}((E_t[ln(1 + \frac{\varphi_t+\xi_{t+1}}{R_{t,t+1}^{1/\rho}}] + u_{t+1})^2)$ depends on the household's level of net assets, and it cannot be neglected since it is not small around small consumption innovations u_{t+1}.

To express u_{t+1} as a function of the income innovations, they approximate each side of the intertemporal budget constraint, which yields the first relation below. Using their assimilation between e and u, they additionally obtain the second relation below:

$$\begin{cases} u_{t+1} + \mathcal{O}(E_t[u_{t+1}^2]) & = \varepsilon_{t+1}\phi_t^{\varepsilon} + \eta_{t+1}\phi_t^{\eta} + \pi_t + \mathcal{O}((\varepsilon_{t+1}, \eta_{t+1})^2) \\ \mathcal{O}(E_t[u_{t+1}^2]) & = \mathcal{O}(\varepsilon_{t+1}^2, \eta_{t+1}^2). \end{cases}$$

This finally yield that log-consumption growth is linear in income innovations around small income innovations:

$$\Delta ln(c_{t+1}^{BPP}) = \underbrace{ln(R_{t,t+1}^{1/\rho}) + \pi_t}_{\Gamma_{t+1}} + \varepsilon_{t+1}\phi_t^{\varepsilon} + \eta_{t+1}\phi_t^{\eta} + \mathcal{O}((\varepsilon_{t+1}, \eta_{t+1})^2)$$

Appendix C Other differences between the robust estimator and the BPP estimator

Bias from using moments conditionally on the period Instead of pooling all periods at the same time, it is possible to estimate the elasticity of consumption at each period using the covariance and variance conditionally on the period at which they are observed:

$$\hat{\phi}_t^\varepsilon = \frac{cov_t(\Delta ln(c_t), -\Delta ln(y_{t+2}))}{cov_t(\Delta ln(y_t), -\Delta ln(y_{t+2}))} = \phi_t^\varepsilon$$

Yet, because the number of observations at each period is too small, a number of existing semi-structural estimators additionally impose that the elasticity be constant over time: $\phi_t^\varepsilon = \phi^\varepsilon$ for all t. If in fact the elasticity varies over time, in an estimation method that gives more weight to the moments that yields estimates with a small variance, this gives more weight to the years in which the elasticity is small and less weight to the years in which it is large. The estimate obtained is smaller than the average elasticity of consumption over the sample period.

Bias from using biannual data with annual shocks Starting in 1999, the PSID data has been recorded every other year and a number of studies implement the method developed by BPP into the more recent biennial panel data. When income shocks are actually annual, it is no longer possible to exactly identify the elasticity of consumption to a contemporaneous transitory shock. Indeed, log-consumption growth and log-income growth are now computed over two periods:

$$\Delta^2 ln(c_t) = \phi^\eta \eta_t + \phi^{\eta L1}\eta_{t-1} + \phi^\varepsilon \varepsilon_t + (\phi^\varepsilon + \phi^{\varepsilon L1})\varepsilon_{t-1} + (\phi^{\varepsilon L1} + \phi^{\varepsilon L2})\varepsilon_{t-2} + (\phi^{\varepsilon L2} + \phi^{\varepsilon L3})\varepsilon_{t-3}$$
$$\Delta^2 ln(y_t) = \eta_t + \eta_{t-1} + \varepsilon_t + \theta\varepsilon_{t-1} - \varepsilon_{t-2} - \theta\varepsilon_{t-3}$$

The ratio of the covariances of log-consumption growth and log-income growth with future log-income growth writes:

$$\begin{aligned}\hat{\phi}^\varepsilon &= \frac{cov_t(\Delta^2 ln(c_t), -\Delta^2 ln(y_{t+2}))}{cov_t(\Delta^2 ln(y_t), -\Delta^2 ln(y_t))} \\ &= \frac{cov_t(ln(c_t), \varepsilon_t + \theta\varepsilon_{t-1})}{var_t(\varepsilon_t + \theta\varepsilon_{t-1})} \\ &= \phi^\varepsilon + \underbrace{(1-\theta)\phi^\varepsilon \frac{\theta var_t(\varepsilon_{t-1})}{var_t(\varepsilon_t + \theta\varepsilon_{t-1})}}_{(1)>0} + \underbrace{\phi^{\varepsilon L1} \frac{\theta var_t(\varepsilon_{t-1})}{var_t(\varepsilon_t + \theta\varepsilon_{t-1})}}_{(2)<0}.\end{aligned}$$

The elasticity of interest is still the elasticity of consumption to a transitory shock over the period observed, that is ϕ^ε. On the one hand, even if consumption growth was uncorrelated with past shocks, this ratio of covariance would not measure our coefficient of interest because of the transitory shock at $t-1$ fading away from period t and not from period $t+1$ on. This leads to an underestimation of the variance of $\varepsilon_t + \theta\varepsilon_{t-1}$, thus to an overestimation of the true elasticity. I denote this effect (1). On the other hand, because consumption growth is strictly negatively correlated with the realization of past shocks, the positive consumption effect of a positive transitory shock at $t+1$ is mitigated by its negative consumption effect at t and it is not possible to disentangle the two. This leads to an underestimation of the covariance between consumption growth and $\varepsilon_t + \theta\varepsilon_{t-1}$, thus to an underestimation of the true elasticity. I denote this effect (2).

Bias from measurement error in combination with a non-zero MA coefficient θ Neglecting the presence of classical measurement error can also lead one to use additional moments of the variance-autocovariance of log-income growth that do not hold when there is measurement error. The effect of measurement error depends on the persistence of the transitory income process. Indeed, although the covariance between current log-income growth and future log-income growth two periods later identifies $\theta var(\varepsilon_t)$, which is sufficient to get a consistent estimator, it is tempting to use other moments of the log-income growth process to get more precise estimates of the variance of the shocks and therefore of the consumption elasticity. Besides the use of the covariance between log-income growth and log-income growth two periods later, the other non-zero moments are the covariance between log-income growth and log-income growth one period later, and the variance of log-income growth. Yet these additional moments come with additional parameters to estimate. In particular, to help measuring the variance of the transitory shock more precisely, they require some external identification of the variance of measurement error and of the MA coefficient θ, which otherwise cannot be disentangled from the variance of the transitory shock. Although the identification problem is clearly stated in Meghir and Pistaferri (2004), who acknowledge that 'it is not possible to disentangle the unconditional variance of the transitory shock, the variance of the measurement error, and the MA coefficients.'(p.11), it is overlooked in the estimation developed by BPP: the authors assume there is no measurement error in the baseline case and discuss its effect in the Appendix (Appendix C of the full web Appendix), in the case with an MA(0) transitory income process. In that case, neglecting measurement error leads to an overestimation of the variance of the transitory shock and an underestimation of the elasticity of consumption to a transitory shock.

I add two elements to this discussion. First, there is an additional identification problem associated with the MA coefficient, related to the fact that it is the solution of a quadratic equation. It means that the system of equation should give two possible roots for θ in general, and not just one value. Second, I show that, in the presence of measurement error, using the covariance between log-income growth and log-income growth one period later as an estimating moment in addition to the covariance between log-income growth and log-income growth two periods later leads to overestimating the highest root of θ and underestimating the smallest root of θ. The estimate of the MA coefficient obtained by BPP is smaller than one, which I show corresponds to the smallest root, so it is biased downward if measurement error is important. Because the product $\theta var(\varepsilon_t)$ is correctly identified, an underestimation of θ translates into an overestimation of $var(\varepsilon_t)$. This eventually generates a bias of undetermined direction in the estimation of the elasticity of consumption to a transitory shock that depends on the additional moments used.

References

Abowd, J. and Card, D. (1989). "On the Covariance Structure of Earnings and Hours Changes." *Econometrica* 57.2, pp. 411–45.

Agarwal, S., Liu, C., and Souleles, N. (2007). "The Reaction of Consumer Spending and Debt to Tax Rebates-Evidence from Consumer Credit Data." *Journal of Political Economy* 115.6, pp. 986–1019.

Agarwal, S and Qian, W. (2014). "Consumption and Debt Response to Unanticipated Income Shocks: Evidence from a Natural Experiment in Singapore." *American Economic Review* 104.12, pp. 4205–30.

Arrow, K. (1965). "Aspects of the theory of risk-bearing." Yrj Jahnssonin Sti (Yrjo Jahnsson Lectures).

Attanasio, O. and Davis, S. (1996). "Relative Wage Movements and the Distribution of Consumption." *Journal of Political Economy* 104.6, pp. 1227–62.

Attanasio, Orazio, Meghir, Costas, and Mommaerts, Corina (2015). "Insurance in extended family networks." mimeo.21059. URL: `https://ideas.repec.org/p/nbr/nberwo/21059.html`.

Baker, S. and Yannelis, C. (2015). "Income Changes and Consumption: Evidence from the 2013 Federal Government Shutdown." mimeo.

Bayer, C. and Juessen, F. (2015). "Happiness and the Persistence of Income Shocks." *American Economic Journal: Macroeconomics* 7.4, pp. 160–87. DOI: `10.1257/mac.20120163`.

Blundell, R., Graber, M., and Mogstad, M. (2016). "Labor income dynamics and the insurance from taxes, transfers, and the family." mimeo.

Blundell, R., Low, H., and Preston, I. (2013). "Decomposing changes in income risk using consumption data." *Quantitative Economics* 4.1, pp. 1–37.

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

Blundell, R., Pistaferri, L., and Saporta-Eksten, I. (2016). "Consumption inequality and family labor supply." *American Economic Review* 106.2, pp. 387–435.

Brav, A., Constantinides, G., and Geczy, C. (2002). "Asset Pricing with Heterogeneous Consumers and Limited Participation: Empirical Evidence." *Journal of Political Economy* 110.4, pp. 793–824.

Carlos Hatchondo, J., Martinez, L., and Sànchez, J. (2015). "Mortgage defaults." *Journal of Monetary Economics* 76.C, pp. 173–190. DOI: `10.1016/j.jmoneco.2015.09`. URL: `https://ideas.repec.org/a/eee/moneco/v76y2015icp173-190.html`.

Casado, J. (2011). "From income to consumption: measuring households partial insurance." *Empirical Economics* 40.2, pp. 471–495. DOI: `10.1007/s00181-010-0337-z`.

Cochrane, J. (1991). "A Simple Test of Consumption Insurance." *Journal of Political Economy* 99.5, pp. 957–76.

Commault, J. (2016). "Precautionary Saving and the Heterogeneity of the Consumption Response to Shocks." mimeo.

Dynarski, S. and Gruber, J. (1997). "Can Families Smooth Variable Earnings?" *Brookings Papers on Economic Activity* 28.1, pp. 229–303.

Etheridge, B. (2015). "A test of the household income process using consumption and wealth data." *European Economic Review* 78.C, pp. 129–157. DOI: `10.1016/j.euroecorev.2015`.

Fagereng, A., Holm, M., and Natvik, G. (2016). "MPC Heterogeneity and Household Balance Sheets." mimeo.

Gelman, M., Kariv, S., Shapiro, M., Silverman, D., and Tadelis, S. (2015). "How Individuals Smooth Spending: Evidence from the 2013 Government Shutdown Using Account Data." mimeo.

Heathcote, J., Storesletten, K., and Violante, G. (2014). "Consumption and Labor Supply with Partial Insurance: An Analytical Framework." *American Economic Review* 104.7, pp. 2075–2126.

Hedlund, A., Karahan, F., Mitman, K., and Ozkan, S. (2017). *Monetary Policy, Heterogeneity and the Housing Channel*. Tech. rep. 663.

Hryshko, D. (2014). "Correlated income shocks and excess smoothness of consumption." *Journal of Economic Dynamics and Control* 48.C, pp. 41–62. DOI: 10.1016/j.jedc.2014.08.02.

Johnson, D., Parker, J., and Souleles, N. (2006). "Household Expenditure and the Income Tax Rebates of 2001." *American Economic Review* 96.5, pp. 1589–1610.

Kan, K., Peng, S, and Wang, P (2016). "Understanding Consumption Behavior: Evidence from Consumers Reaction to Shopping Vouchers." mimeo.

Kaplan, G. and Violante, G. (2010). "How Much Consumption Insurance beyond Self-Insurance?" *American Economic Journal: Macroeconomics* 2.4, pp. 53–87.

Kaplan, G. and Violante, G. (2014). "A Model of the Consumption Response to Fiscal Stimulus Payments." *Econometrica* 82.4, pp. 1199–1239.

Kaufmann, K. and Pistaferri, L. (2009). "Disentangling Insurance and Information in Intertemporal Consumption Choices." *American Economic Review* 99.2, pp. 387–92. DOI: 10.1257/aer.99.2.387.

Kimball, M. (1990). "Precautionary Saving in the Small and in the Large." *Econometrica* 58.1, pp. 53–73.

MaCurdy, T. (1982). "The use of time series processes to model the error structure of earnings in a longitudinal data analysis." *Journal of Econometrics* 18.1, pp. 83–114.

Misra, K. and Surico, P. (2014). "Consumption, Income Changes, and Heterogeneity: Evidence from Two Fiscal Stimulus Programs." *American Economic Journal: Macroeconomics* 6.4, pp. 84–106.

Ortigueira, Salvador and Siassi, Nawid (2013). "How important is intra-household risk sharing for savings and labor supply?" *Journal of Monetary Economics* 60.6, pp. 650–666. DOI: 10.1016/j.jmoneco.2013.05.

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

Pistaferri, L., Saporta-Eksten, I., and Blundell, R. (2017). "Children, Time Allocation and Consumption Insurance." mimeo.

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

Souleles, N. (1999). "The Response of Household Consumption to Income Tax Refunds." *American Economic Review* 89.4, pp. 947–958.

Zheng, Y. and Santaeulalia, R. (2016). "The Price of Growth: Consumption Insurance in China 1989-2009." mimeo.826.