How Does Consumption Respond to a Transitory Income Shock? Departing From a Random Walk Consumption to Reconcile Semi-Structural Estimates With Natural Experiment Results

Jeanne Commault*

February 2018

(Click here for the most recent version)

Abstract

I show that a generalized version of the semi-structural estimation method developed by Blundell, Pistaferri, and Preston (2008), which relaxes their assumption that log-consumption is a random walk, estimates the elasticity of consumption to a transitory income shock to be large. With the same PSID dataset, the average yearly elasticity of nondurable consumption is 0.59, a magnitude consistent with findings from natural experiments. I explain that the seminal random walk expression of consumption initially developed by Hall (1978) is not an approximation of the standard life-cycle model but erroneously obtained by linearizing an identity. Contrary to this, in general, the standard life-cycle model generates a consumption process that departs from a random walk because of precautionary behavior. Numerical simulations show that, when calibrated with the persistence of the transitory income process and of the variance of the transitory and permanent shocks that I re-estimated, the standard life-cycle model generates an elasticity of 0.42.

Key words: Marginal propensity to consume, transitory income shocks, random walk hypothesis

*Department of Economics, Sciences Po, 28 Rue des Saints-Pères, 75007 Paris, France; jeanne.commault@sciencespo.fr. I thank Richard Blundell, Stéphane Bonhomme, Pierre Cahuc, Edouard Challe, Mariacristina De Nardi, Juan Dolado, Stéphane Dupraz, Jean-Marc Robin, Ben Etheridge, Francis Kramarz, Hamish Low, Eric Mengus, Xavier Ragot, Michèle Tertilt, and participants at seminars where the paper was presented for helpful comments.

Introduction

How does individual consumption respond to a transitory income shock? An important obstacle in the way of measuring this response 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 the difference is important to a researcher, 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 a tax rebate and a lottery win, and pairing them with consumption data to directly measure the response of expenditures to an income shock that the researcher observes and knows to be transitory. In the great majority of these studies, a transitory income change has a statistically significant and economically large effect on consumption.[1] A second approach identifies the response of consumption to a transitory shock 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 a transitory and to a permanent shock. Specifically, the authors assume that income evolves as a transitory-permanent process, and that log-consumption is a random walk so that past shocks affect past and current log-consumption exactly in the same way and have no effect on log-consumption growth.[2] This identification strategy is now influential in all fields of economics that are concerned with the way shocks are passed on to households' decision variables such as consumption and saving but also individual labor supply and time allocation. Yet, the studies relying on this estimation method typically find that the elasticity of consumption to a transitory shock is not statistically significant, although it is quite precisely estimated.[3] Examining the validity of the semi-structural estimator in life-cycle models, Kaplan and Violante (2010)

[1] See for example Parker (1999), Souleles (2002) for a transitory change in take-home pay; 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 refund or tax rebate; Baker and Yannelis (2017), Gelman, Kariv, Shapiro, Silverman, and Tadelis (2018) for the response to the 2013 government shutdown; Agarwal and Qian (2014), Kan, Peng, and Wang (2017) for the response to the distribution of cash or consumption vouchers by the government; Fagereng, Holm, and Natvik (2018) for the response to a lottery win. A related but distinct literature relies on hypothetical survey responses rather than direct observations of consumption to measure how households respond to a transitory shock. See Parker and Souleles (2017) for a comparison between hypothetical survey measures and natural experiment measures.

[2] This method is said to be semi-structural because it does not require that the fully-fledged life-cycle model holds but only that this one condition derived from an approximation of this model does, together with the transitory-permanent specification of income.

[3] See section A of the appendix for a review of the literature that builds on the BPP identification strategy. Early papers before BPP study the response of consumption to a change in total income (Krueger and Perri (2005) Krueger and Perri (2011)), 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)).

conclude that the absence of correlation between log-consumption growth and past shocks, which they conjecture might not hold exactly in a life-cycle model, is not even required for the identification of the elasticity of consumption to a transitory shock.

In this paper, I make four contributions: (i) I show that, although the Kaplan and Violante (2010) version of the BPP estimator of the elasticity of consumption to a transitory shock is immune to the presence of a correlation with past shocks, the original BPP method is not; (ii) I implement in survey data a robust version of the BPP estimator that allows log-consumption growth to depend on past income shocks, and I find that the elasticity to a contemporaneous transitory shock becomes large and statistically significant; (iii) I investigate the theoretical conditions under which log-consumption correlates with past shocks: I show that precautionary behavior generates a departure from a random walk, which does not disappear in first order approximation around small innovations to income or consumption; I explain that the seminal random walk expression of Hall (1978), the log-linearized Euler equation, and other previous approximations neglect the correlation because they are not approximating the first-order condition of the household problem but an identity; (iii) I examine the quantitative importance of these effects in numerical simulations. A life-cycle model calibrated using the income process parameters that I have re-estimated without bias make it possible to produce an elasticity close to my empirical estimate.

Thus, firstly, the general identification strategy of the BPP estimator is to use future log-income growth as an instrument to identify the effect of the transitory shock separately from that of the permanent shock. Indeed, the current transitory shock covaries negatively with future log-income growth, as its effect fades away in the periods after it realizes, but the current permanent shock does not affect future log-income growth, as it raises current and future log-income by the same amount. Using only future log-income growth at the last period in the future before the effect of a current transitory shock fully dissipates, the instrument correlates with the current transitory shock but not with the past transitory shocks: it is robust to the presence of a correlation between log-consumption growth and past shocks. The original BPP method, however, uses all observations of future log-income growth as instruments, assuming away the effect of past transitory shocks on current log-consumption growth. Thus, if the realizations of past transitory shocks correlate in fact with current log-consumption growth, the instrument is endogenous and the estimator biased. In addition, the original BPP method uses other moments that are not robust in the estimation, although these moments likely identify other parameters than the elasticity and would not induce much further bias in the measure of this elasticity. Among related methods, the one developed by Arellano, Blundell, and Bonhomme (2017), which considers a more general income process, is not robust to any types of correlation but more robust than the original BPP in that it lets log-consumption growth depend on current assets and current permanent income, thus allowing it to correlate with past shocks through the effect of those past shocks on current assets and permanent income.[4] The technique

[4]Yet, as they focus on permanent shocks, they only estimate the elasticity to a transitory shock on simulated

developed by Blundell, Pistaferri, and Saporta-Eksten (2016) to further estimate Frisch elasticities from the BPP estimates is biased in the presence of a correlation between log-consumption growth and past shocks.

Secondly, I implement the robust version of the BPP estimator in data from the Panel Study of Income Dynamics (PSID) between 1978 and 1992, combined with consumption data imputed from the Consumer Expenditure Survey (CEX) over the same period. This follows the original BPP study, except that I additionally detrend from the effect of past demographic characteristics because serial correlation in characteristics could otherwise shift the estimate. With the robust estimator, the average yearly elasticity of nondurable consumption to a transitory shock has a point estimate of 0.59 and is statistically significant. An estimator that is identical except for including future log-income growth at periods when it correlates with past transitory shocks in the set of instruments yields a point estimate of 0.02, not statistically significant. The full original BPP estimator, which uses future income growth at periods when it correlates with past transitory shocks as instrument and additionally relies on other moments to jointly estimate extra parameters, obtains a similar point estimate of 0.01.[5] The average yearly marginal propensity to consume (MPC) out of current transitory income that the elasticity implies is at least 0.32, which is consistent with results from natural experiments. I further estimate the change in elasticity over time. I obtain a point estimate of -0.52, meaning the elasticity of consumption to a transitory shock over the second year after that shock hit is smaller than the elasticity over the first year by 0.52. I also re-estimate the income parameters, and find the persistence of the transitory process to be larger, the variance of the transitory shocks smaller, and the variance of the permanent shocks larger. The robust estimator of the elasticity to a transitory shock is by construction robust to the presence of classical measurement error. It is also robust to letting the shocks be continuous instead of discrete (which is a possible bias noted in Crawley (2018)). Allowing households to anticipate the shocks to the extent that their anticipations are consistent with the covariances between log-consumption and future log-income, or allowing for serial correlation in measurement error can only raise the estimate of the elasticity to a transitory shock. Finally, I find that modeling permanent income as an AR(1) with a coefficient smaller than one only raises the elasticity estimate as well.

Thirdly, I confirm analytically the conjecture of Kaplan and Violante (2010) that in the standard life-cycle model log-consumption growth is not a random walk and can correlate with past shocks. The source of the correlation is precautionary behavior, that is, the effect of uncertainty on the behavior of a household. The correlation induced by precautionary behavior does not disappear in first order approximation around small fluctuations in the realization of future consumption or income, because precautionary saving is decided before future consumption and income is realized and depends on the possibility that it departs from its expected value and not on the realized departure from its expected value. The reason why the correlation with

data and not on survey data.

[5] It reaches 0.05 when the data is not detrended from the effect of past characteristics, as in the original paper.

past shocks is nevertheless commonly approximated away is that, since Hall (1978), random walk expressions of log-consumption growth are based on the linearization of an identity rather than on the linearization of the first order condition of the household maximization problem. Also, the exact random walk expression of consumption derived by Hall (1978) under quadratic preferences does not require that consumption be positive. Imposing such a constraint makes the solution for consumption sensitive to uncertainty even under quadratic preferences, and generates a departure from a random walk.

Fourthly, I measure the quantitative importance of these effects in numerical simulations that mimic the PSID 1978-1992 dataset. I simulate a standard life-cycle model, the same model as Kaplan and Violante (2010) that they design to broadly reproduce the BPP data, in which (i) I let transitory income be an MA(1) (ii) I use the persistence of transitory income and variances of the transitory and permanent shocks that I re-estimated, (iii) instead of having two polar constraints (zero borrowing or only the natural borrowing constraint), I consider the intermediate case in which households can borrow up to a fraction of their income. The model generates an elasticity of consumption to a transitory income shock of 0.42.

1 Estimating the elasticity of consumption

1.1 Statistical model

The statistical model that I assume encompasses the standard life-cycle model as a special case. In particular, it does not impose that a household solves a maximization problem to choose its consumption. The estimating restrictions are as follows.

Log-income growth Log-income is a transitory-permanent income process, the sum of a permanent component that evolves as a random walk and of a transitory component that evolves as an $MA(q)$. The order q of the MA process is established empirically: in the dataset I consider, the covariance between current log-income growth and future log-income growth is no longer statistically significant after $t+2$, so that $q=1$ and I denote $\theta_1=\theta$. Log-income can also depend linearly on demographic characteristics and other shocks $\zeta^y_{i,t}$ that may capture measurement error. It implies that the log-income growth of household i at period t detrended from the linear effect of its demographic characteristics z, denoted $\Delta ln(\tilde{y}_{i,t})$, is a linear function of the current permanent shock, the current transitory shock, and the past transitory shocks up to period $t-q-1=t-2$:

$$\Delta ln(\tilde{y}_{i,t}) = \eta_{i,t} + \varepsilon_{i,t} - (1-\theta)\varepsilon_{i,t-1} - \theta\varepsilon_{i,t-2} + \zeta^y_{i,t} - \zeta^y_{i,t-1}. \tag{1.1}$$

Log-consumption growth The log-consumption growth of a household i at period t detrended from the linear effect of its demographic characteristics z, denoted $\Delta ln(\tilde{c}_{i,t})$, is a linear function of the current and past transitory shocks ε it has received, and a flexible function of the current and past permanent income shocks η it has received and on other consumption-specific shocks ζ^c:

$$\Delta ln(\tilde{c}_{i,t}) = \phi^{\varepsilon}_{i,t}\varepsilon_t + (\phi^{\varepsilon L1}_{i,t} - \phi^{\varepsilon}_{i,t-1})\varepsilon_{t-1} + ... + (\phi^{\varepsilon L(t-1)}_{i,t} - \phi^{\varepsilon L(t)}_{i,t-1})\varepsilon_1 \\ + f(\eta_{i,t}, ..., \eta_{i,1}, \zeta^c_{i,t}, ..., \zeta^c_{i,1}). \quad (1.2)$$

The function $f(.)$ relating permanent income shocks and consumption-specific shocks to consumption can be non-linear. As in the case of income, the shocks ζ^c can represent measurement error. Note that, instead of assuming that log-consumption growth is linear in the current and past transitory shocks, it is alternatively possible to assume that transitory shocks are normally distributed (though the permanent shocks can be distributed in anyway).

Distributional assumptions I make the following assumptions about the distributions of the variables:

· the demographic characteristics of a household affect its log-consumption and log-income growth linearly; there is no need for the demographic characteristics of a household to be independent from the shocks it receives: if they do correlate, the elasticity I measure is the direct effect of the shock on consumption, excluding the indirect effect through its interaction with demographic characteristics.;

· the different shocks received by a household are drawn independently of one another and independently over time;

For the standard error to be consistently estimated, I additionally assume that shocks are drawn independently between households, though I allow for arbitrary within-household correlations. There is no assumption imposing uniformity in the distributions from which households draw their shocks, since the estimation procedure is robust to heteroskedasticity: different households can draw their shocks from different distributions, and the same household can draw its shocks from different distributions over time.

Household information The model does not impose that households know about their income process, nor that they can distinguish between the permanent and transitory shocks they receive. If they do not, they will simply respond in the same way to a permanent shock as they do to a transitory shock.

Average elasticity Under these distributional assumptions, the ratio of the covariance between log-consumption growth and a transitory shock over the variance of this shock coincides with the average elasticity of consumption to a transitory shock in the population, denoted ϕ^{ε}:

$$\hat{\phi}^{\varepsilon} = \frac{cov(\Delta ln(\tilde{c}_{i,t}), \varepsilon_{i,t})}{var(\varepsilon_{i,t})} = E\left[\frac{d\Delta ln(\tilde{c}_{i,t})}{d\varepsilon_{i,t}}\right] = E[\phi^{\varepsilon}_{i,t}]. \tag{1.3}$$

1.2 Identification: instrumenting with future income growth

When the realizations of the shocks ε are observed, typically in the context of natural experiments, it is straightforward to measure the covariance $cov(\Delta ln(\tilde{c}_{i,t}), \varepsilon_{i,t})$ and variance $var(\varepsilon_{i,t})$, and to estimate the elasticity to a transitory shock. In survey data, however, the realizations of the shocks ε are not directly accessible. Only total income y is reported, and current log-income growth is driven by the realizations of several different shocks: the current permanent shock, the current transitory shock, and the past transitory shocks.

The solution is to instrument the effect of current log-income growth with future log-income growth, which correlates with the realization of the current transitory shock but not with the realization of the current permanent shock. Indeed, a positive transitory income shock raises log-income at t, increasing log-income growth from $t-1$ to t, but as its effect dissipates it does not raise log-income by as much at $t+1$ and no longer raises it at $t+2$, decreasing log-income growth from t to $t+1$ and from $t+1$ to $t+2$. Contrary to this, as a positive permanent shock raises log-income once and for all, so it increases log-income growth from $t-1$ to t but does not decrease it at subsequent periods. In addition, log-income growth at $t+2$ correlates with the realization of the current transitory shock but not with the realizations of any past transitory shocks, whose effects are already fully dissipated. Thus, the covariance between log-consumption growth at t and log-income growth at $t+2$ is exclusively driven by the realization of the transitory shock at t:

$$cov(\Delta ln(\tilde{c}_{i,t}), -\Delta ln(\tilde{y}_{i,t+2})) = \theta cov(\Delta ln(\tilde{c}_{i,t}), \varepsilon_{i,t}). \tag{1.4}$$

Similarly, the covariance between log-income growth at t and log-income growth at $t+2$ is driven by the realization of the transitory shock at t:

$$cov(\Delta ln(\tilde{y}_{i,t}), -\Delta ln(\tilde{y}_{i,t+2})) = \theta var(\varepsilon_{i,t}). \tag{1.5}$$

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

$$\hat{\phi}^{\varepsilon} = \frac{cov(\Delta ln(\tilde{c}_{i,t}), -\Delta ln(\tilde{y}_{i,t+2}))}{cov(\Delta ln(\tilde{y}_{i,t}), -\Delta ln(\tilde{y}_{i,t+2}))} = \phi^{\varepsilon}. \tag{1.6}$$

1.3 Bias from neglecting the correlation with past shocks

Using endogenous instruments Neglecting the presence of a correlation between log-consumption growth and the realizations of past shocks leads a researcher to use additional instruments that are endogenous when such a correlation is present. Indeed, when past income shocks have no effect on current log-consumption growth, it is possible to use also the opposite of future log-income growth at $t+1$ as an instrument, to get more identifying moments. When past shocks do have an effect, log-income growth at $t+1$ correlates with log-consumption growth through both the current transitory shock and the immediately past transitory shock, so that using it as an additional instrument without acknowledging the effect of the past shock on log-consumption growth yields a bias. In particular, when a past transitory shock correlates negatively with log-consumption growth, because it correlates positively with the opposite of future log-income growth, it reduces the covariance between log-consumption growth and the additional instrument:

$$\begin{cases} cov(\Delta ln(\tilde{c}_{i,t}), -\Delta ln(\tilde{y}_{i,t+2})) &= \theta cov(\Delta ln(\tilde{c}_{i,t}), \varepsilon_{i,t}), \\ cov(\Delta ln(\tilde{c}_{i,t}), -\Delta ln(\tilde{y}_{i,t+1})) &= (1-\theta) cov(\Delta ln(\tilde{c}_{i,t}), \varepsilon_{i,t}) + \theta \underbrace{cov(\Delta ln(\tilde{c}_{i,t}), \varepsilon_{i,t-1})}_{\text{neglected}}. \end{cases} \tag{1.7}$$

The underestimation of this covariance induces an underestimation of the elasticity of consumption to a transitory shock. The two expressions simultaneously identifying the elasticity consumption to a transitory shock are:

$$\begin{cases} \hat{\phi}^{\varepsilon}_{BPP} &= \frac{cov(\Delta ln(\tilde{c}_{i,t}), -\Delta ln(\tilde{y}_{i,t+2}))}{cov(\Delta ln(\tilde{y}_{i,t}), -\Delta ln(\tilde{y}_{i,t+2}))} = \phi^{\varepsilon}, \\ \hat{\phi}^{\varepsilon}_{BPP} &= \frac{\theta}{1-\theta} \frac{cov(\Delta ln(\tilde{c}_{i,t}), -\Delta ln(y_{i,t+1}))}{cov_t(\Delta ln(\tilde{y}_{i,t}), -\Delta ln(\tilde{y}_{i,t+2}))} = \phi^{\varepsilon} + \frac{\theta}{1-\theta} \underbrace{\frac{cov(\Delta ln(\tilde{c}_{i,t}), \varepsilon_{i,t-1})}{var_t(\varepsilon_{i,t})}}_{\text{neglected}} \neq \phi^{\varepsilon}. \end{cases} \tag{1.8}$$

The first identifying expression is unbiased, it is the same as in the robust estimator, but the second expression underestimates the elasticity. The negative effect of the past transitory shock on the covariance between log-consumption growth and the additional instrument is erroneously attributed to the fact that log consumption responds less to the current transitory shock than it really does.

Additionally using an erroneous estimate of the persistence of a transitory shock In addition, using log-income growth at $t+1$ as an instrument requires having an estimate of the persistence θ. This means that an additional bias $(\hat{\theta}-\theta)cov(\Delta ln(\tilde{c}_{i,t}), \varepsilon_{i,t})$ can arises in the measure of the elasticity if the estimate of θ used is not its true value:

$$cov(\Delta ln(\tilde{c}_{i,t}), -\Delta ln(\tilde{y}_{i,t+1})) = (1-\hat{\theta})cov(\Delta ln(\tilde{c}_{i,t}), \varepsilon_{i,t}) + \underbrace{\theta cov(\Delta ln(\tilde{c}_{i,t}), \varepsilon_{i,t-1})}_{\text{neglected}} + \underbrace{(\hat{\theta}-\theta)cov(\Delta ln(\tilde{c}_{i,t}), \varepsilon_{i,t})}_{\text{neglected}}. \tag{1.9}$$

Although the larger the true value of θ is, the larger the bias from using an endogenous instrument, the effect is reversed for the estimated value of θ: the lower below the true value is the estimate, the larger the bias from misestimating θ (provided that the direction of the covariance with a past shock is the opposite as that of the covariance with a current shock). In the original BPP, the estimation of θ is based on moments that neglect the presence of measurement error in income, so it si likely subject to some bias.

Other sources of bias in the original BPP estimator The original BPP method is subject to the two biases listed above, as it relies on future log-income growth at $t+1$ as an instrument and the value of theta on which it relies is estimated. In addition, the method simultaneously estimates a host of parameters along with the elasticity of consumption to a transitory shock, from additional moments. In particular, the variance of the transitory shocks is estimated at each period and identified from additional moments besides (1.5), which also omit the presence of a correlation with past shocks and the presence of measurement error in income. This can indirectly affect the estimation of the elasticity of consumption to a transitory income shock. Note that although this elasticity also enters the covariance between current log-consumption and log-income growth, which is used in the BPP specification, this moment is the only one that identifies the elasticity of consumption to a permanent shock so that any misspecification in this moment would show up in the estimate of the latter and not of the former.

Bias in related estimation methods I show in Section B of the appendix, that the presence of a correlation with past shocks induces a similar bias in the estimator of the Frisch elasticities of consumption developed by Blundell, Pistaferri, and Saporta-Eksten (2016).

1.4 Change in elasticity over time

The covariance between current log-consumption growth and future log-income growth makes it possible to identify the covariance between current log-consumption growth and the past transitory shock, conditionally on simultaneously estimating the covariance between log-consumption and the current shock and the persistence θ:

$$cov(\Delta ln(\tilde{c}_{i,t}), -\Delta ln(\tilde{y}_{i,t+1})) = (1-\theta)cov(\Delta ln(\tilde{c}_{i,t}), \varepsilon_{i,t}) + \theta cov(\Delta ln(\tilde{c}_{i,t}), \varepsilon_{i,t-1}). \tag{1.10}$$

Combined with (1.4) and (1.5), I can further identify the change in elasticity over time:

$$(\widehat{\phi^{\varepsilon}-\phi^{\varepsilon L1}}) = \frac{cov(\Delta ln(\tilde{c}_{i,t}), -\Delta ln(\tilde{y}_{i,t+1}))}{cov(\Delta ln(\tilde{y}_{i,t-1}), -\Delta ln(\tilde{y}_{i,t+1}))} - \frac{1-\theta}{\theta}\frac{cov(\Delta ln(\tilde{c}_{i,t}), -\Delta ln(\tilde{y}_{i,t+2}))}{cov(\Delta ln(\tilde{y}_{i,t-1}), -\Delta ln(\tilde{y}_{i,t+1}))}. \quad (1.11)$$

This persistence θ can be jointly estimated from the covariance between current and future log-income growth, conditionally on simultaneously estimating the variance of the current and past transitory shock and on knowing the value of the variance of measurement error in income:

$$cov(\Delta ln(\tilde{y}_{i,t}), -\Delta ln(\tilde{y}_{i,t+1})) = (1-\theta)var(\varepsilon_{i,t}) + \theta var(\varepsilon_{i,t-1}) + var(\zeta^{y}_{i,t}). \quad (1.12)$$

The point that the persistence cannot be separately identified from measurement error was first made in Meghir and Pistaferri (2004). They solve is by using an external estimate of measurement error close to 0.01, which is what I do as well.

2 Results

2.1 Data and estimator

Data I implement the robust estimator in data from the Panel Study of Income Dynamics (PSID) between 1978 and 1992, which contains longitudinal information on a representative sample of US households surveyed every year. This PSID data is combined with consumption data imputed from the Consumer Expenditure Survey (CEX) over the same period. I detail the sample selection and the definitions of variables in Appendix C.

Detrending difference with BPP: removing the effect of past demographic characteristics The only difference with the data used in BPP is my including past demographic characteristics in the set of detrending variables. Indeed, the underlying assumption of the detrending strategy used in the original BPP method is that log-consumption and log-income depend only on the current demographic characteristics. The rationale for detrending is then to avoid the correlation between the demographic characteristics influencing consumption and income showing up as a response of consumption to the income shocks. As the robust estimator relies only on covariances with future log-income growth at $t+2$, it is less subject to this problem and would only be biased in the presence of serial correlation in the demographic characteristics if some of these characteristics were unobserved. However, if past demographic characteristics also influence current consumption and income (because of frictions or adjustment periods), correlations between demographic characteristics arise again in the estimating moments, including correlations between variables observed at the same period.[6] For this reason, I break

[6]When the demographic characteristics are independent of the realizations of the current and past shocks that are captured in the residuals, the only difference between using detrended and non-detrended variables in the model is the presence of the terms $cov(\Delta\delta_t z_{i,t}, \Delta\kappa_{t+2} z_{i,t+2})$ and $cov(\Delta\kappa_t z_{i,t}, \Delta\kappa_{t+2} z_{i,t+2})$, which are non-zero only

away from the choices made in the original BPP method and additionally include the value of the same demographics characteristics at the past period. I do not include variables at $t+2$, because it does no longer make a difference, as shown in Appendix F, but reduces the number of observations thus the precision. I also include interactions between those demographics and education dummies, as done later in Blundell, Pistaferri, and Saporta-Eksten (2016).

Estimator I implement the estimators described by (1.6) and by the combination of (1.6) and (1.8) in this dataset with a generalized method of moment that is detailed in Appendix D.

2.2 Estimating moments

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

Covariances	$\Delta ln(\tilde{y}_{i,t})$	$\Delta ln(\tilde{y}_{i,t+1})$	$\Delta ln(\tilde{y}_{i,t+2})$	$\Delta ln(\tilde{y}_{i,t+3})$
$cov(\Delta ln(\tilde{c}_{i,t}),.)$	0.0097	0.0017	-0.0041	0.0000
	(0.0017)	(0.0016)	(0.0017)	(0.0015)
$cov(\Delta ln(\tilde{y}_{i,t}),.)$	0.0679	-0.0194	-0.0069	0.0009
	(0.0028)	(0.0014)	(0.0013)	(0.0014)
Obs.	7,600	7,600	7,600	6,285

Note: Standard errors in parentheses are adjusted for arbitrary within-household correlations and heteroskedasticity.

Covariances of log-consumption growth Before looking at the elasticity estimates, Table 1 presents the value of the moments used for estimation. The first line shows the covariances between log-consumption growth and current and future log-income growth. They are consistent with my statistical model, but inconsistent with a model that does not allow for a correlation between log-consumption growth and past income shocks, because the covariance between log-consumption growth and future log-income growth at $t+1$ is positive. In the absence of a correlation with past shocks it should be be negative and correspond to a fraction $(1-\theta)/\theta$ of the covariance with future log-income growth at $t+2$ (from equations (1.7)). In the presence of a correlation with past shocks, however, an additional term is present in this moment, $\theta cov(\Delta ln(\tilde{c}_{i,t}), \varepsilon_{i,t+1})$, which can overturn its sign.

Covariances of log-income growth The autocovariances of log-income growth presented in the second line of Table 1 suggest that 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

in the presence of serial correlation in z. When past demographics influence current consumption and income, the difference is the presence of the terms $cov(\Delta(\delta_t z_{i,t} + \delta_t z_{i,t-1}), \Delta(\kappa_{t+2} z_{i,t+2} + \kappa_{t+2} z_{i,t+1}))$ and $cov(\Delta(\kappa_t z_{i,t} + \kappa_t z_{i,t-1}), \Delta(\kappa_{t+2} z_{i,t+2} + \kappa_{t+2} z_{i,t+1}))$, which are non zero even in the absence of serial correlation.

log-income growth 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 log-income growth at $t+3$ should be statistically different from zero, while it is not. Also, if permanent income was not a random walk but an AR(1) with a coefficient different from one, the autocovariances between log-income growth at t and at all future periods should be statistically different from zero, while they stop being significant after two periods. In case the AR(1) coefficient is just slightly below one, generating a covariance with future log-income growth that is too small to be precisely estimated, I consider AR(1) models in the extensions.

2.3 Elasticity to a transitory shock

Table 2: Elasticity ϕ^ε

	Robust	Non-robust	BPP	BPP
ϕ^ε	0.594 (0.250)	0.023 (0.050)	0.012 (0.046)	0.053 (0.043)
θ	*n.a.* *n.a.*	0.203 *n.a.*	0.203 (0.024)	0.113 (0.025)
$\underline{MPC}^\varepsilon$	0.319 (0.135)	0.012 (0.027)	0.006 (0.025)	0.028 (0.023)
Obs.	7,600	7,600	9,052	10,472
Add. dem.	yes	yes	yes	no
Moments	(1.6)	(1.6), (1.8)	$(1.6)_t$, $(1.8)_t$ $\forall t$ + others	$(1.6)_t$, $(1.8)_t$ $\forall t$ + others

Note: Standard errors in parentheses are adjusted for arbitrary within-household correlations and heteroskedasticity. The number of observations is the number of household-year pairs for which (1.6) or (1.10) are observed (in the second column I keep only observations for which both are observed so the sample is the same as in the first column).

Robust estimator The first column of Table 2 reports the results obtained with my robust estimator, based on moment (1.6), which holds even in the presence of a correlation between log-consumption growth and past income shocks. Using this estimator, the average elasticity of nondurable consumption to a transitory shock on net income is large, with a point estimate of 0.59, statistically significant at 5%. It means that a transitory shock that raises current income by 10% and next period income by $\theta \times 10\%$ leads to a 5.9% increase in current nondurable consumption over the following year, on average in the sample.

Non-robust estimators The next three columns show 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 (1.10), an expression that neglects a term that is non-zero in the presence of such a correlation. These estimators yield much smaller estimates of the average elasticity of consumption to a transitory income shock than the robust estimator. The second column features the results from a simple non-robust estimator that is identical to the robust estimator except that it additionally relies on moment (1.10). The associated estimate of the elasticity of consumption to a transitory shock is 0.01, much below the estimate of 0.59 obtained with the robust estimator, and not statistically significant. The fact that the point estimate is substantially smaller is consistent with my theoretical prediction that using moment (1.10) for estimation induces a downward bias in the measure of the elasticity of consumption to a transitory income shock when log-consumption growth and past transitory shocks correlate negatively. The third column presents the results obtained with the original BPP estimator, which uses moments (1.6) and (1.10) as well, but differs from the robust estimator and the simple non-robust estimator on other grounds because it relies on additional moments and considers all moments conditionally on the period. The point estimate of the average elasticity of consumption to a transitory income shock is similar to that of the simple non-robust estimator, at 0.01, and not statistically significant. The similarity in results between the third and fourth columns suggests that, although additional estimating moments are used in the BPP estimator, they serve to identify other parameters and the elasticity to a transitory shock remains identified mainly from moments (1.6) and (1.10). Finally, the fourth column presents the results obtained with the original BPP estimator and data detrended with the same, more limited, set of demographics considered in the original paper.

Marginal Propensity to Consume (MPC^{ε}) Semi-structural estimators measure the elasticity of consumption to a transitory shock $\phi^{\varepsilon}_{i,t} = \frac{1}{c_{i,t}} \frac{dc_{i,t}}{d\varepsilon_{i,t}}$, that is, the percentage change in current consumption associated with a one unit transitory income shock, which causes a 100% change in current income and a $\theta \times 100\%$ change in future income.[7] On the contrary, natural experiments typically measure the MPC out of current income, that is, the level change in current consumption caused by a level change in current income: $MPC^{\varepsilon}_{i,t} = \frac{dc_{i,t}}{dy_{i,t}} = \frac{d\varepsilon_{i,t}}{dy_{i,t}} \frac{dc_{i,t}}{d\varepsilon_{i,t}} = \frac{1}{y_{i,t}} \frac{dc_{i,t}}{d\varepsilon_{i,t}}$. The relation between the elasticity of consumption and the MPC out of the change in current income caused by a current transitory shock is:

$$MPC^{\varepsilon}_{i,t} = \frac{c_{i,t}}{y_{i,t}} \phi^{\varepsilon}_{i,t}.$$

A difficulty is that I do not measure the individual elasticities $\phi^{\varepsilon}_{i,t}$ but only the average elasticity in the sample $\phi^{\varepsilon} = E[\phi^{\varepsilon}_{i,t}]$. Under the assumption that the individual elasticities are all equal (and thus equal to the average elasticity), the average MPC is the product of the average elasticity and the average ratio of consumption over income. Under the alternative assumption that

[7] The change in current income caused by a transitory shock is $\frac{dy_{i,t}}{d\varepsilon_{i,t}} = \frac{de^{p_{i,t}} e^{\varepsilon_{i,t}} e^{\theta \varepsilon_{i,t-1}}}{d\varepsilon_{i,t}} = y_{i,t}$: a one unit increase in $\varepsilon_{i,t}$ raises current income by 100%, that is, doubles it. The change in future income caused by a transitory shock is $\frac{dy_{i,t+1}}{d\varepsilon_{i,t}} = \frac{de^{p_{i,t+1}} e^{\varepsilon_{i,t+1}} e^{\theta \varepsilon_{i,t}}}{d\varepsilon_{i,t}} = \theta y_{i,t+1}$: a one unit increase in $\varepsilon_{i,t}$ raises current income by $\theta \times 100\%$.

individual elasticities are not necessarily all equal, but are such that the households with the highest elasticities are on average those with the highest ratios of consumption over income, this product is a lower bound for the average MPC:

$$\underline{MPC}^{\varepsilon} = E[\frac{c_{i,t}}{y_{i,t}}]\phi^{\varepsilon} \leq E[\frac{c_{i,t}}{y_{i,t}}\phi^{\varepsilon}_{i,t}] = MPC^{\varepsilon}.$$

Because the average ratio of consumption over income is smaller than one, the lower bound for the average MPC is smaller than the average elasticity. I find that a household consumes on average at least 45% of the change in its current income caused by a transitory shock. The lower bounds for the average MPC measured with non-robust estimators are small, at 4%, and not statistically significant.

Comparison with the literature on natural experiments How do this elasticity and these MPCs compare with the results derived from natural experiments? Studying increases in take-home pay, Parker (1999) finds that the average elasticity of nondurable consumption over the next three months out of a temporary increase in take-home pay caused by a change in social security taxes is 0.54, significant at 1%. Souleles (2002) estimates the MPC of nondurable consumption over the next three months out of a change in take-home pay induced by the Reagan tax cuts to be 0.66, significant at 5%. Baker and Yannelis (2017) and Gelman, Kariv, Shapiro, Silverman, and Tadelis (2018) estimate the MPC of nondurable consumption and of total credit card spending over the next two weeks out of the temporary decrease in take-home pay caused by the 2013 government shutdown to be 0.39 and 0.58, both statistically significant at 1%. Studies of tax refunds, tax rebates, and stimulus program receipts, that are mailed directly to the households, obtain MPCs of nondurable expenditures over the next three months that are between 0.09 and 0.37.[8] Kan, Peng, and Wang (2017) consider the effect of a shopping voucher program, and find that it stimulates reported total expenses above what would have been spent otherwise by 0.24 of the voucher value over the next three months, significant at 1%. Fagereng, Holm, and Natvik (2018) find that the MPC of total spending (measured in a broad way as the difference between income growth and wealth growth) out of a small lottery prize (below $2,150) over the next year is 1.01, significant at 1%. Thus, two characteristics seem robust to the idiosyncrasies of the shocks considered in these different natural experiments: (i) the MPC of nondurable consumption out of an unexpected transitory shock is statistically significant (ii) its value over the year following the shock is at least 0.10. The MPC derived from a robust estimator is consistent with both stylized facts, while the MPCs derived from non-robust estimators conflict with both.

Robustness to continuous arrival of shocks, anticipations, serial correlation in measure-

[8] See Souleles (1999), Johnson, Parker, and Souleles (2006), Parker, Souleles, Johnson, and McClelland (2013), Misra and Surico (2014), and Agarwal and Qian (2014).

ment error, MA(1) permanent income process I explain in Appendix E that the robust estimator is unaltered by shocks being a continuous process rather than discretely occurring every year. It is conservative if shocks can be partly anticipated by the households, for the patterns of anticipation that are consistent with the data. It is also conservative if measurement error is not classical but presents some serial correlation. Finally, allowing log-income to be an MA(1) process rather than a random-walk induces a bias of undetermined direction. Yet, conditionally on knowing the MA(1) coefficient, a consistent estimator can be obtained. I implement a version of the robust estimator that allows for an MA(1) permanent income with a MA(1) coefficient comprised between 1 (baseline) and 0.91. I obtain higher estimated elasticities as my assumption about the MA(1) coefficient moves below 1.

Alternative specifications I check that the findings are robust to the choice of demographic characteristics, interactions, clusters, and measures of consumption and income. Results are reported in section F of the appendix. Incidentally, I find that the elasticity to a transitory shock remains large and significant provided that variables are detrended from the effect of past demographics and employment status, which is the variable most likely to be serially correlated. This is consistent with my model, in which the sources of bias are the influence of past characteristics and the presence of serial correlation in demographic characteristics. The results for food expenditures, total expenditures (nondurables and durables), and total expenditures plus expenditures on education and health, which can be considered an investment, are consistent with comparable natural experiments findings.The response of consumption to shocks on measures of income that exclude taxes or exclude both taxes and transfers is smaller, in line with the role that taxes and transfers are expected to play: it means that, when households receive a shock on their income before taxes and transfers, they anticipate adjustments in taxes and transfers that will reduce the magnitude of the shock, so they respond less to it than they do to a shock on income after taxes and transfers.

2.4 Change in elasticity over time

Table 3: Elasticity ϕ^{ε}

	Robust
ϕ^{ε}	0.622
	(0.249)
$(\phi^{\varepsilon L1} - \phi^{\varepsilon})$	-0.546
	(0.395)
θ	0.797
	(0.187)
$var(\zeta^{y}_{i,t})$	0.010
	n.a.
$\underline{MPC}^{\varepsilon\ total}$	0.176
	(0.074)
Obs.	6,337
Add. dem.	yes
Moments used	(1.6), (1.8), (1.9)

Note: see the main text.

Joint estimation of the elasticity, the change in elasticity, and the persistence Table 3 presents the results of a joint estimation of the elasticity of consumption to a current transitory shock, than change in elasticity over time, and the persistence θ, conditional on making an assumption about the variance of the measurement error on log-income. As it requires four consecutive observation of log-income growth instead of three when estimating solely the elasticity to a current shock, the sample is reduced to 6,337 observations. On this reduced sample, the elasticity of consumption to a current transitory shock is 0.622, very similar to its value of 0.593 on the baseline sample. The change from the elasticity to a current transitory shock to the elasticity to a past transitory shock is -0.546. Although it is not precisely estimated, this means that the difference is potentially large, and explains why neglecting the term $(\phi^{\varepsilon L1} - \phi^{\varepsilon})\varepsilon_{i,t-1}$ in log-consumption growth can generate a substantial bias. This point estimate would mean that the elasticity of consumption to a past transitory shock is only $0.622 - 0.546 = 0.076$: a past transitory shock that multiplied past log-income by 10% and multiplies current log-income by 7.97% only multiplies current consumption by 0.76%. The effect of a transitory shock on consumption would thus drop sharply, consistent with results from natural experiments. Finally, the persistence of a transitory income shock is large and precisely estimated. It is much higher than the estimate obtained with the original BPP method, with a point estimate of 0.797

instead of 0.203 when the original BPP method applied to similarly detrended data. This means that the original BPP is not only using an endogenous instrument but also an underestimated persistence, which amplifies the bias from the endogenous instrument.

Marginal Propensity to Consume out of Total Income Change ($MPC^{\varepsilon\, total}$) While the transitory shocks identified in survey data are found to evolve as MA(1), affecting both current and future income, most of the shocks considered in natural experiments are purely transitory, affecting current income only. Now that I have an estimate of the persistence θ, I can attempt at taking this difference into account. Considering the limit case in which a change in future income has no effect on current consumption, for instance because of borrowing constraints or myopic behavior, this difference in persistence is not a problem, and the MPC out of the change in current income caused by an MA(1) transitory shock also measures how households would respond to a purely transitory shock. Considering the opposite case in which a change in future income has exactly the same effect on current consumption as a change in current income, it is the MPC out of the total net present value change in income caused by an MA(1) transitory shock that measures how households would respond to a purely transitory shock: $MPC^{\varepsilon\, total}_{i,t} = \frac{dc_{i,t}}{dy_{i,t}} \frac{1}{y_{i,t}+\theta/(1+r)y_{i,t+1}}$.[9] The relation with the elasticity is:

$$MPC^{\varepsilon\, total}_{i,t} = \frac{c_{i,t}}{y_{i,t} + (\theta/(1+r))y_{i,t+1}} \phi^{\varepsilon}_{i,t}.$$

Measuring this MPC requires taking a stand on the value of the interest factor $(1+r)$ and on the persistence of the transitory shock θ. I select $(1+r) = 1.03$ and I use the value of $\theta = 0.797$ that I estimated. A lower bound for the average MPC out the total net present value change in income caused by a transitory shock is:

$$MPC^{\varepsilon\, total} = E[\frac{c_{i,t}}{y_{i,t} + \theta/(1+r)y_{i,t+1}} \phi^{\varepsilon}_{i,t}] \geq E[\frac{c_{i,t}}{y_{i,t} + \theta/(1+r)y_{i,t+1}}]\phi^{\varepsilon} = \underline{MPC}^{\varepsilon\, total}.$$

With the large persistence that I have estimated, the total income change caused by a transitory shock is substantially larger than just the change in current income. The change in consumption relative to this total change is smaller . A lower bound for $MPC^{\varepsilon\, total}$ is 0.176, so a household consumes at least 18% of the total net present value gain in income caused by a transitory shock over the following year.

[9] I compute the MPC out of the realized change in current and future income caused by a current transitory shock, and not the MPC out of the expected change $y_{i,t} + (\theta/(1+r))E[y_{i,t+1}]$, because I do not observe expected income $E[y_{i,t+1}] = e^{p_t} E_t[e^{\eta_{i,t+1}}]E_t[e^{\varepsilon_{i,t+1}}]$.

3 Theoretical Correlation With Past Shocks and Elasticity to a Transitory Shock

3.1 The Standard Life-cycle Model

Income process Time is discrete and indexed by t = 0, 1, ... , T. The net income of a household i at period t, denoted $y_{i,t}$, is modeled as a transitory-permanent process:

$$ln(y_{i,t}) = p_{i,t} + \mu_{i,t} + f_i + \kappa_t z_{i,t} \tag{2.1}$$

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

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 $f_i + \kappa_t z_{i,t}$ that captures individual fixed effects and the deterministic influence of current demographic characteristics $z_{i,t}$. The uncertainty of the household about its future income comes from the presence of the shocks, $\eta_{i,t}$ 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 the income received by the household for the rest of its lifetime. The shock $\varepsilon_{i,t}$ is transitory because its effect on income dissipates after q periods.[10] At each period, the permanent and transitory shocks are drawn independently of each other and independently of their previous realizations. The demographic characteristics $z_{i,t}$ are not subject to any uncertainty: they may change over time, but these variations are expected by the household. Their impact on log-income is measured by the vector of coefficients κ_t, which is allowed to change with calendar time.

Household's problem At each period t, a household i chooses its current consumption and the distribution of its future consumption as the solution of the following intertemporal optimization problem:

$$\max_{c_{i,t},\ldots c_{i,T}} \quad \sum_{s=0}^{T-t} \beta^{t+s} e^{\delta_{t+s} z_{i,t+s}} E_t \left[u(c_{i,t+s}) \right] \tag{2.2}$$

$$s.t. \quad a_{i,t+k+1} = (1+r) a_{i,t+k} - c_{i,t+k} + y_{i,t+k} \qquad \forall\, 0 \leq k \leq T-t, \tag{2.3}$$

$$a_{i,T} \geq 0. \tag{2.4}$$

The household is finite-lived with T the length of its life. It has time-separable preferences, and at each period t it derives utility from its contemporaneous 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 by the factor β, and shifted by the demographic characteristics $z_{i,t}$, whose

[10] By construction, at the end of the household's life the transitory shock resembles a permanent shock (whose effect is depreciating), because its effect might last until the last period of the household's life.

current and future values are known in advance with certainty by the household. The impact of these demographics on utility is captured by the vector of coefficients δ_t, which can change with calendar time. At each period, the household earns the stochastic amount $y_{i,t}$. The budget constraint (2.3) states that to store its wealth from one period to another the household only has access to a risk-free asset that delivers the certain interest rate r, where $a_{i,t}$ denotes the level of this asset at the beginning of period t or at the end of period $t-1$. The terminal condition on wealth (2.4) states that the household cannot die with a strictly positive level of debt.

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

$$u'(c_{i,t}) = E_t[u'(c_{i,t+1})]R_{i,t,t+1},$$

with $R_{i,t,t+1} = \beta(1+r)e^{\Delta\delta_{t+1}z_{i,t+1}}$ a factor accounting for the deterministic intertemporal substitution motives. It states that an optimizing household chooses its current and future consumption so that they deliver the same expected marginal utility. The natural borrowing limit never binds, as the household would never put itself in the situation of possibly consuming zero in the future by borrowing more than the lowest possible amount that it could earn in the future. The effect of intertemporal substitution can equivalently be expressed as a weight $R_{i,t,t+1}^{1/\rho}$ on current consumption: $u'(c_{i,t})R_{i,t,t+1}^{-1} = c_{i,t}^{-\rho}R_{i,t,t+1}^{-1} = (c_{i,t}R_{i,t,t+1}^{1/\rho})^{-\rho}$.

Contribution of precautionary behavior to consumption growth When marginal utility is convex, the effects of negative and positive shocks are asymmetric: a negative shock to future consumption raises the value of one additional unit of future consumption more than a positive shock of the same magnitude reduces the value of one additional unit of future consumption. On average, the effect of the negative shocks dominates, so the presence of mean-zero shocks increases the expected marginal utility of future consumption above the marginal utility of expected future consumption: $E_t[u'(c_{i,t+1})] > u'(E_t[c_{i,t+1}])$. This induces the household to set its current consumption below its expected future consumption, so the marginal utility of its current consumption, which is not subject to uncertainty, be as high as the expected marginal utility of its future consumption, which is increased by the uncertainty about future consumption. The amount by which the household must decrease its current consumption coincides with the risk-premium associated with the marginal utility of consumption, that is, the variable $\varphi_{i,t}$ such that: $E_t[u'(c_{i,t+1})] = u'(E_t[c_{i,t+1}] - \varphi_{i,t})$, which is strictly positive when marginal utility is convex. I refer to $\varphi_{i,t}$ as precautionary consumption growth because it corresponds to the effect of uncertainty on consumption growth.[11]

$$\begin{aligned} u'(c_{i,t}R_{i,t,t+1}^{1/\rho}) &= u'(E_t[c_{i,t+1}] - \varphi_{i,t}), \\ E_t[c_{i,t+1}] &= c_{i,t}R_{i,t,t+1}^{1/\rho} + \underbrace{\varphi_{i,t}}_{\text{precaution}} . \end{aligned}$$

Preservation through first order approximations The contribution of precaution to expected consumption growth is unaffected by approximations around small shocks to realized future income, or to realized future consumption. This is because the term arises from the possibility of such shocks, and does not depend on whether shocks effectively realized or not.

3.2 Pitfalls in previous approximations of consumption growth

Hall (1978)'s random walk expression under isoelastic utility Hall also shows that, when utility is not quadratic but isoelastic, consumption still approximately evolves as a random walk around small shocks. I explain that this expression is not derived from approximating the first order condition of the household's problem, but from approximating an identity plugged in this first order condition. More precisely, what Hall (1978) does is starting from the first order condition, substituting for $E_t[u'(c_{i,t+1})] = u'(c_{i,t+1}) + (E_t[u'(c_{i,t+1})] - u'(c_{i,t+1}))$, and applying $(u')^{-1}(.)$ to each side:

$$\begin{aligned} u'(c_{i,t}R_{i,t,t+1}^{1/\rho}) &= E_t[u'(c_{i,t+1})] \\ u'(c_{i,t}R_{i,t,t+1}^{1/\rho}) &= u'(c_{i,t+1}) + (E_t[u'(c_{i,t+1})] - u'(c_{i,t+1})) \\ c_{i,t+1} &= (u')^{-1}(u'(c_{i,t+1}) + \cancel{E_t[u'(c_{i,t+1})]} - \cancel{u'(c_{i,t}R_{i,t,t+1}^{1/\rho})}). \end{aligned}$$

The crossing of the terms that cancel out is mine. It shows that the first order condition is unessential here, and that the expression of $c_{i,t+1}$ obtained is the same as the identity $c_{i,t+1} = (u')^{-1}(u'(c_{i,t+1}))$. Indeed, because the first order condition does not establish a relation between current consumption and realized future consumption but only between current consumption and the expected distribution of future consumption, trying to express realized future consumption from it leads to relying on an identity. Hall (1978) then takes a first order approximation of $c_{i,t+1}$ around the point where $u'(c_{i,t+1}) = u'(c_{i,t}R_{i,t,t+1}^{1/\rho})$, i.e. uses $f(x) \approx$

[11] Kimball (1990) relies on the risk-premium as well, but uses it to prove the slightly different result that, in a tractable two-period model, a household facing uncertainty must be holding strictly more assets as it would in the absence of uncertainty to chose the same level of consumption. I make the point that it also coincides with the difference in consumption in the absence and presence of uncertainty, thus coincides with the difference between the current, certain, consumption of a household and its future, uncertain, consumption that equalizes their marginal utility.

$f(x_0) + (x - x_0)f'(x_0)$ with $f(x) = c_{i,t+1}$, $x = u'(c_{i,t+1})$, and $x_0 = u'(c_{i,t}R^{1/\rho}_{i,t,t+1})$:

$$c_{i,t+1} \approx c_{i,t}R^{1/\rho}_{i,t,t+1} + \underbrace{\frac{u'(c_{i,t+1}) - u'(c_{i,t}R^{1/\rho}_{i,t,t+1})}{u''(c_{i,t}R^{1/\rho}_{i,t,t+1})}}_{\text{uncorrelated with past shocks}}.$$

This is Hall (1978)'s random walk expression, in which future consumption is the sum of current consumption plus a mean zero innovation, uncorrelated with past shocks because $u'(c_{i,t+1}) - u'(c_{i,t}R^{1/\rho}_{i,t,t+1}) = u'(c_{i,t+1}) - E_t[u'(c_{i,t+1})]$. Yet, this expression reflects the choice of the identity that is linearized: if instead one were to approximate $c_{i,t+1} = v^{-1}(v(c_{i,t+1}))$, with $v(.)$ such that $E_t[v(c_{i,t+1})] \neq v(c_{i,t}R^{1/\rho}_{i,t,t+1})$, at the same point as Hall (1978) where $c_{i,t+1} = c_t R^{1/\rho}_{i,t,t+1}$, the resulting expression would not be a random walk, although the first order approximation would hold.[12] The particular identity $c_{t+1} = (u')^{-1}(u'(c_{i,t+1}))$ yields a random walk because it relies on a choice of $x = u'(c_{i,t+1})$ such that $E_t[x] = x_0$.

Although it is not what Hall (1978) does, it would be possible to derive a random walk expression of consumption from the first order condition of the household maximization problem by taking an approximation of current consumption around the point where the expected distribution of future consumption is a Dirac delta function with mass one in current consumption, that is, where future consumption equals current consumption in all states of the world, regardless of the shocks that realize.[13] This is substantially more restrictive than an approximation around the point where realized future consumption equals current consumption, which can occur in the standard life-cycle model with uncertainty, while the point where future consumption equals current consumption in all states of the world cannot emerge as an outcome of the model.

Also, the exact random walk expression of consumption does not hold when consumption is constrained to be positive, because a natural borrowing constraint emerges that can then be binding in the presence of uncertainty. Intuitively, such a constraint is equivalent to creating a kink at zero that renders marginal utility convex. In the presence of uncertainty about future consumption, and when this uncertainty is such that future consumption is not distributed only on the linear part of the marginal utility function, it generates precautionary consumption growth. Current consumption is not equal to future expected consumption, but additionally

[12]The resulting expression would be $c_{i,t+1} \approx c_{i,t}R^{1/\rho}_{i,t,t+1} + \frac{v(c_{i,t+1}) - v(c_{i,t}R^{1/\rho}_{i,t,t+1})}{v''(c_{i,t}R^{1/\rho}_{i,t,t+1})}$ with $v(c_{i,t}R^{1/\rho}_{i,t,t+1}) \neq E_t[v(c_{i,t+1})]$.

[13]In the discrete case with J possible future states of the world, and each state $j \in J$ occurring with probability π_j:

$$u'(c_{i,t}R^{1/\rho}_{i,t,t+1}) = \sum_{j\in J} \pi_j u'(c^j_{i,t+1}) \approx \sum_{j\in J} \pi_j u'(c_{i,t}R^{1/\rho}_{i,t,t+1}) + \sum_{j\in J} \pi_j (c^j_{i,t+1} - c_{i,t}R^{1/\rho}_{i,t,t+1}) u''(c_{i,t}R^{1/\rho}_{i,t,t+1})$$
$$0 = E_t[c_{i,t+1}] - c_{i,t}R^{1/\rho}_{i,t,t+1}.$$

incorporates a precautionary term.

The log-linearized Euler equation The log-linearized Euler equation is typically obtained with same method as Hall (1978)'s, by linearizing the identity $ln(c_{i,t+1}) = ln(c_{i,t}R^{1/\rho}_{i,t,t+1}) - \frac{1}{\rho}ln(1+\frac{u'(c_{i,t+1})-u'(c_{i,t}R^{1/\rho}_{i,t,t+1})}{u'(c_{i,t}R^{1/\rho}_{i,t,t+1})})$ around $u'(c_{i,t+1}) = u'(c_{i,t}R^{1/\rho}_{i,t,t+1})$. The second order version of the log-linearized Euler equation does not express log-consumption growth as a random walk but comes from the approximation of an identity as well (and all higher-order versions of the log-linearized Euler equation do).

BPP's approximation The expression used by BPP is based on the one derived by Blundell, Low, and Preston (2013), which partly follows a similar procedure as Hall (1978)'s, applied to log-consumption.

3.3 Correlation with past transitory shocks

Theorem: In the model presented above, there exist values of the variance of the transitory and permanent shocks such that the precautionary consumption growth φ_t is strictly negatively correlated with current assets and with the value of a past transitory income shock. At any period $0 < t < T$:

$$\exists\ \sigma^{\eta}_{i,t+1}\ s.t.\ \frac{d\varphi_{i,t}}{da_{i,t}} < 0 \text{ and } \frac{d\varphi_{i,t}}{d\varepsilon_{i,t-k}} < 0,\quad k \geq q.$$

Proof: First, I show that precautionary consumption growth is decreasing in assets. Indeed, although it is known since Leland (1968) and Sandmo (1970) that precautionary saving is decreasing in assets,[14] it does not directly follow that precautionary consumption growth should be decreasing in assets as well. The latter result requires that a gain in assets reduces current precautionary saving more than it reduces expected future precautionary saving. To prove this result, I derive both sides of the first order condition of the maximization problem with respect

[14]More precisely, see Leland (1968), Sandmo (1970), Drèze and Modigliani (1972),Sibley (1975), and Miller (1976), for the derivation and progressive generalization of this result. Kimball (1990), often cited on the subject, proves in fact the slightly different result that to chose the same level of consumption as in the absence of uncertainty a household facing uncertainty must be holding strictly more assets.

to a change in assets, and divide by $(-u''(c_t R^{1/\rho}_{i,t,t+1}))$:

$$\frac{dc_{i,t}R^{1/\rho}_{i,t,t+1}}{da_{i,t}} = E_t[\frac{dc_{i,t+1}}{da_{i,t}} \frac{u''(c_{i,t+1})}{u''(c_{i,t}R^{1/\rho}_{i,t,t+1})}]$$

$$\frac{dc_{i,t}R^{1/\rho}_{i,t,t+1}}{da_{i,t}} = \frac{dE_t[c_{i,t+1}]}{da_{i,t}} \underbrace{\frac{E_t[-u''(c_{i,t+1})]}{-u''(c_{i,t}R^{1/\rho}_{i,t,t+1})}}_{>1} + \underbrace{cov_t\left(\frac{dc_{i,t+1}}{da_{i,t}}, \frac{-u''(c_{i,t+1})}{-u''(c_{i,t}R^{1/\rho}_{i,t,t+1})}\right)}_{\exists\ \sigma^{\varepsilon}_{i,t+1}\ s.t.\ >0}.$$

Two effects lead a household to reduce its precautionary consumption growth in response to a gain in assets, under certain conditions on the variance of the permanent shock. The first effect is that a gain in assets shifts the expected distribution of future consumption upwards, to a region where the convexity of marginal utility is less pronounced. More precisely, when marginal utility is convex, the change in marginal utility $-u''(.)$ is a convex function of marginal utility $u'(.)$. It means that the value of current consumption that equalize the current and future expected marginal utility do not equalize their changes: $E_t[-u''(c_{i,t+1})] > -u''(c_{i,t})$. Then, although current consumption is lower than expected consumption so that increasing current and expected consumption by one unit reduces the marginal utility of current consumption more $-u''(c_{i,t}) > -u''(E_t[c_{i,t+1}])$, increasing current and future consumption by one unit reduces the uncertainty premium to expected marginal utility more: $E_t[-u''(c_{i,t+1})] - (-u''(E_t[c_{i,t+1}])) > (-u''(c_{i,t})) - (-u''(E_t[c_{i,t+1}]))$. Overall, one additional unit of future consumption reduces the expected marginal utility of future consumption more than one additional unit of current consumption reduces the marginal utility of current consumption, and a household optimally responds by raising its current consumption more than its future expected consumption, reducing its precautionary consumption growth.

The second effect comes from the fact that a change in assets raises future consumption in all states of the world but not by the same amount $\frac{dE_t[c_{i,t+1}]}{da_{i,t}}$ in all states. In Commault (2018b), I complete the proof of Carroll and Kimball (1996) that consumption is concave in assets and in transitory income. Concavity means that having more assets reduces one's response to a transitory income shock $\frac{dc_{i,t+1}}{da_{i,t+1}\varepsilon_{i,t+1}} < 0$. Having more assets at t, and therefore more assets at $t+1$, raises consumption more in the states of the world in which low transitory shocks realize, thus in which consumption will be low everything else being equal. This reduces the variance in future consumption caused by transitory income shocks. Yet this might not be true of permanent shocks: a gain in current assets might not raise future consumption more in the states of the world in which low permanent shocks realize, so a gain in current assets might not reduce the variance of future consumption caused by the realizations of future permanent shocks. When the variance of the permanent shocks is small enough, though, the reduced uncertainty coming from the response to transitory shocks dominates. Then, future consumption responds more to a change in current assets in the states of the world when future consumption is low, that is, in which the sensitivity of marginal utility to consumption is high: $cov_t\left(\frac{dc_{i,t+1}}{da_{i,t}}, \frac{-u''(c_{i,t+1})}{-u''(c_{i,t}R_{i,t,t+1}^{1/\rho})}\right) > 0$.[15] This positive covariance means that the expected marginal utility of future consumption falls even more with a gain in assets, and a household must raise its current consumption even more than its future expected consumption for its current and expected marginal utility to fall by the same amount, reducing even more its precautionary consumption growth.

Second, I show that having experienced a positive transitory shock in the past raises strictly the current level of assets of a household. Deriving each side of the first order condition with respect to a transitory income shock and rearranging yields: $\frac{dc_t}{d\varepsilon_t} = (1+r)\frac{E_t[\frac{dc_{t+1}}{da_{t+1}}\frac{-u''(c_{t+1})}{-u''(c_t)}]}{1+E_t[\frac{dc_{t+1}}{da_{t+1}}\frac{-u''(c_{t+1})}{-u''(c_t)}]} < (1+r)$. Indeed, a household experiencing a positive transitory shock never consumes it entirely in one period, because this would imply that its current marginal utility would decrease while its future expected marginal utility would not, so the two would no longer be equal. Since part of the transitory income gain is passed on to the next period, it generates an increase in future assets. This increase is not entirely consumed either, so that eventually a gain in transitory income increases future assets at all subsequent periods until the end the household's lifetime: $\frac{da_{t+s}}{d\varepsilon_t} > 0$ for all $s > 0$. This means that the realization of any past transitory shock is positively correlated with current assets. For all $0 < s < t$:

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

Implications for the correlation with log-consumption growth As precautionary behavior modifies consumption growth, it modifies log-consumption growth as well: the latter incorporates a term that depends positively on precautionary consumption growth φ_t and on the innovation to consumption $\nu_{i,t+1}$, scaled by current consumption $c_{i,t}R_{i,t,t+1}^{1/\rho}$:

$$\Delta ln(c_{i,t+1}) = \underbrace{\frac{1}{\rho}ln(R_{i,t,t+1})}_{\substack{\text{change in dem.}\\ \text{+ int. substitution}\\ \text{(deterministic)}}} + \underbrace{E_t[ln(1+\frac{\varphi_{i,t}+\nu_{i,t+1}}{c_{i,t}R_{i,t,t+1}^{1/\rho}})]}_{\substack{\text{precaution}\\ \text{(correlates with past shocks)}}} + \underbrace{\xi_{i,t+1}}_{\substack{\text{innovation}\\ \text{(uncorrelated}\\ \text{with past shocks)}}}.$$

The second term on the right-hand-side is the contribution of precaution to expected log-consumption growth, because it would be zero under perfect foresight (i.e. no uncertainty

[15] Formally, the covariance rewrites:

$$cov_t\left(\frac{dc_{i,t+1}}{da_{i,t}}, \frac{-u''(c_{i,t+1})}{-u''(c_{i,t}R_{i,t,t+1}^{1/\rho})}\right) = \underbrace{cov_t\left(\frac{dc_{i,t+1}}{da_{i,t}}(E_t[\eta_{t+1}]), \frac{-u''(c_{i,t+1})}{-u''(c_{i,t}R_{i,t,t+1}^{1/\rho})}(E_t[\eta_{t+1}])\right)}_{>0} + g_t(var_t(\eta_{t+1}))$$

with $g_t(var_t(\eta_{i,t+1})) = 0$ when $var_t(\eta_{i,t+1}) = 0$. The first term is strictly positive because conditional on $\eta_{i,t+1} = E_t[\eta_{i,t+1}]$ the only source of uncertainty about $\frac{dc_{i,t+1}}{da_{i,t}}$ and $\frac{-u''(c_{i,t+1})}{-u''(c_{i,t}R_{i,t,t+1}^{1/\rho})}$ at period t is the realization of the transitory shock at $t+1$ and both terms covary positively with the value this shock, so their covariance is positive.

about the future, everything else being equal).[16] Having received a positive transitory shock in the past can reduce $\varphi_{i,t}$, as I establish in the theorem, so it can reduce the contribution of precaution to expected log-consumption growth. This means that, even in the standard life-cycle model, the log-consumption growth of a household can correlate negatively with the value of the transitory shocks it has received in the past through the contribution of precautionary behavior.

3.4 Elasticity to a Current Transitory Shock

I also look into the implications of precautionary behavior for the consumption elasticity ϕ_t^{ε}. To do so, I first express consumption 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:

$$E_t[c_{t+s}] = c_t R_{t,t+s}^{1/\rho} + \underbrace{\sum_{k=1}^{s} E_t[\varphi_{t+k-1}] R_{t+k,t+s}^{1/\rho}}_{\text{precaution}}.$$

What is the additional amount of saving at t necessary to implement this precautionary consumption growth between t and all future periods? [17] I plug the equation above into the intertemporal budget constraint (2.4), to obtain the following equilibrium relationship satisfied by consumption:

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

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. 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[18] Using equation

[16]This term forms the contribution of precaution to *expected* log-consumption growth but is not the sole precautionary component of *total* log-consumption growth, since the distribution of the innovation is affected by precautionary behavior.

[17]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.

[18]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. 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

(2.7), it is then possible to explicit the contribution of precautionary behavior to the elasticity of consumption:

$$\phi_t^{\varepsilon} = \frac{\frac{dW_t}{d\varepsilon_t} - \overbrace{\frac{dPG_t}{d\varepsilon_t}}^{\text{precaution (1)}}}{W_t - \underbrace{PG_t}_{\text{precaution (2)}}} > \phi_t^{\varepsilon\ perfect\ foresight}$$

with $W_t = (1+r)a_t + \sum_{s=0}^{T-t} \frac{E_t[y_{t+s}]}{(1+r)^s}$ the total expected resources at t, and $PG_t = \sum_{s=1}^{T-t} l_{t,s} \frac{E_t[\varphi_{t+s}+\tilde{\lambda}_{t+s}]}{(1+r)^s}$ the sum of future expected precautionary growth at date $t+1$. Precautionary behavior raises the elasticity of consumption to a transitory shock above the elasticity that would be obtained under perfect foresight, that is if future income was certain everything else being equal. First, precautionary saving decreases with a transitory shock, reducing the response of consumption to the shock. The numerator of the elasticity of consumption is larger than it would under perfect foresight. Second, a household facing uncertainty makes makes precautionary saving and 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: the denominator of the elasticity of consumption is smaller than it would under perfect foresight.

4 Numerical Simulations

4.1 Description of the set-up

Model The outline and calibration of the model that I simulate follow the one Kaplan and Violante (2010), which I extend to allow for an MA(1) transitory income process. The model corresponds to an enriched version of the standard life-cycle model presented above. Households enter their working life at age 25, work for 35 years, until age 60, and retire for up to 35 years, dying with certainty at age 90 if they are still alive. The probability to die becomes non-zero after retirement. Households seek to maximize the sum of their expected marginal utility of consumption, weighted by a discount factor, under a budget constraint, a positivity constraint on consumption, and possibly a borrowing constraint. During working life, the log of net income is the product of a deterministic experience profile and a transitory-permanent process, the sum of a permanent component that evolves as a random walk and of a transitory component that evolves as an MA(1). After retirement, net income is certain and is a function of lifetime average individual gross income.

the weight put on consumption between the beginning of period t and the beginning of period $t+s+1$.

Calibration I use the same survival probabilities as Kaplan and Violante (2010), from the National Center for Health Statistics. The utility function can be quadratic or isoleastic with a risk-aversion parameter of 2. The interest rate is $r = 0.03$, and the discount factor is $\beta = 1/(1+r) \approx 0.97$. The deterministic experience profile of income is calibrated from the PSID. Shocks are normally distributed. The variance of permanent shocks is set at 0.03. The initial variance of the permanent shocks is set at 0.15 to match the dispersion of household earnings at age 25. The variance of the transitory shock is set at 0.01. The MA(1) coefficient is $\theta = 0.8$, following the estimate I obtained. The social security benefits perceived after retirement are a function of lifetime average individual gross earnings designed to mimic the actual US system: benefits are 90% of average past earnings up to a given bend point, 32% percent from this first bend point to a second bend point, and 15% percent beyond that. These benefits are computed based on gross income, while the income process is calibrated to net income. Gross income is thus obtained by inverting the tax function estimated by Gouveia and Strauss (1994). The baseline utility function is quadratic, since I show that it can also generates precautionary behavior only from the interaction of uncertainty with the natural borrowing constraint (or any stronger constraint). Households can borrow up to a fraction 0.73 of the minimum possible permanent income they can have, which corresponds to $\$20,000$ on average.

Simulation method I simulate an artificial panel of 5,000 households from this model, which yields 170,000 observations of their consumption growth during the 35 years of their working life. I have verified that increasing the sample further does not lead to changes in the results. The model is solved using the method of endogenous grid points developed by **Carroll2006**. The grid for assets has 100 exponentially spaced grid points; the grid for the permanent component has 7 equally spaced points; the grid for the transitory shock has 7 equally spaced points.

4.2 Results

Table 4: Elasticity in a simulated life-cycle model

	ϕ^{ε}	$(\phi^{\varepsilon} - \phi^{\varepsilon L1})$
Baseline	0.424	−0.073
	(0.004)	(0.004)
KV without borrowing constraint	0.065	−0.000
	0.001	0.001
KV with zero borrowing	0.154	−0.068
	0.001	0.001
KV with zero borrowing and QP	0.220	−0.120
	0.001	0.001
KV with zero borrowing, QP, and an MA(1)	0.294	−0.067
	0.002	0.001
KV with zero borrowing, QP, an MA(1), updated variances	0.425	−0.108
	0.004	0.004
Obs.	170,000	170,000

Table 4 shows that the baseline life-cycle model generates a large elasticity of consumption to a transitory shock that is more than two thirds of the empirical estimate and within one standard-deviation of it. It also generates a decrease in the value of the elasticity over time, although the decrease is smaller than the one that I estimate empirically.

The rows below shows how the model of Kaplan and Violante, which initially produces a low elasticity of 0.07, can generate a much larger elasticity when the calibration of the income process, borrowing constraints, and kink in the marginal utility varies. The first two lines are cases that are considered by Kaplan and Violante. In the third row, the preferences are additionally changed from isoelastic to quadratic. Interestingly, the quadratic preferences induces a larger elasticity when uncertainty is sufficiently high. An explanation is that, since quadratic preferences only kink when the possibility of negative consumption arises, they induce households to accumulate much less assets, which raises their response to a transitory shock when their distribution of future consumption actually gets closer to zero consumption. In the fourth row the transitory process additionally shifts from an MA(0) to an MA(1) with persistence 0.08. This substantially increases the elasticity. Finally, in the fifth row the variances of the transitory and permanent shocks move from 0.05 and 0.01 to 0.01 and 0.03, consistent with the values that I re-estimated. This raises the elasticity close to the baseline value produced to the model. It is even a little bit above since the baseline borrowing constraint is more flexible than a zero borrowing condition.

5 Conclusion

In this paper, I show that the standard life-cycle model features a correlation between consumption growth and past transitory shocks that is caused by precautionary behavior, and that more general models possibly incorporate additional sources of correlation. I take stock of this possible correlation and generalize the semi-structural estimator of the elasticity of consumption to a transitory income shock, making it robust to such a correlation. The average elasticity of consumption to a transitory shock becomes statistically significant and its point estimate increases to 0.59, which is ten times larger than with an estimator that relies on instruments that are endogenous to past income shocks. The average marginal propensity to consume out of the total net present value change in income caused by a transitory shock is at least 0.32, consistent with the results obtained in natural experiments of transitory income changes.

What does it imply? First, the consistency of findings between the semi-structural estimation and the natural experiments suggests that the shocks considered in natural experiments are not too different from the typical shocks captured in longitudinal survey data. Thus, the strong response of consumption to a transitory income shock seems to be a widespread phenomena rather than a finding confined to fiscal stimuli and lottery wins. Second, the magnitude of the change in results when shifting from a non-robust to a robust estimator means that the effect of past shocks is not negligible and implies some caution in the use of other non-robust semi-structural techniques, including estimators of the elasticity to a permanent shock, estimators of Frisch elasticities, and estimators implemented in biennial datasets.

Appendix A The BPP estimator in the literature

The BPP estimator has been adapted, extended, and put to use in diverse fields, and I provide a few examples for each. 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 rival specifications of income; Bayer and Juessen (2015) apply it to estimate the response of happiness to transitory and permanent income shocks; Ghosh (2016) extends the BPP method to exploit both the second and third moments of log-income and log-consumption growth.

In labor, Ortigueira and Siassi (2013) and Heathcote, Storesletten, and Violante (2014) use the BPP estimates as a benchmark against 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 wage shocks; Blundell, Pistaferri, and Saporta-Eksten (2018) estimate the elasticity of hours spent with children to transitory and permanent wage shocks.

In development, Attanasio, Meghir, and Mommaerts (2015) compare the elasticity of consumption to transitory and permanent income shocks at the village level and at the individual level, to assess the importance of within-village insurance mechanisms. Santaeulàlia-Llopis and Zheng (2018) measure the evolution of the elasticity of consumption to transitory and permanent income shocks during the period of large and sustained GDP growth in China.

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

Appendix B Estimation of Frisch elasticities

Method Building on the BPP estimator, some studies take a step back from net income, considering shocks that can partly be endogenously insured and estimating the Frisch elasticity of consumption to these shocks, holding the marginal utility of wealth constant, to measure the importance of the changes in margins other than wealth caused by the shock. One of the first to do so is Blundell, Pistaferri, and Saporta-Eksten (2016), who consider the effect of a permanent wage shock on consumption, and the extent to which it is insured through adjustments in the labor supply of the household members. The authors estimate the Frisch elasticities by matching empirical estimates of the elasticities of consumption and labor supply to a permanent wage shock, measured with the BPP method, with their theoretical expressions, which depend on the Frisch elasticities, on the elasticities to shocks on net income (after all insurance has taken place), and on observable institutional parameters (response of taxes in particular). For instance, the theoretical elasticity of consumption to a permanent wage shock is:

$$\phi^{\eta\, BPS}_{\text{wage}} = f(\phi^{\eta\, BPS}_{\text{net. inc.}}, F_{i,t}, I_{i,t}),$$

with $\phi^{\eta\, BPS}_{\text{net. inc.}} = \sum_{s=0}^{T-t} \frac{E_t[y_{t+s}]}{((1+r)^s} / ((1+r)a_t + \sum_{s=0}^{T-t} \frac{E_t[y_{t+s}]}{(1+r)^s})$ an approximation of the theoretical elasticity of consumption to a permanent shock on net income, $F_{i,t}$ the vector of Frisch elasticities, $I_{i,t}$ the vector of institutional parameters, and $f(.)$ a functional form derived from the approximated solution of the elasticity to a permanent wage shock in a life-cycle model. The values of $\phi^{\eta\, BPS}_{\text{net. inc}}$ and $I_{i,t}$ are externally measured and plugged in to estimate $F_{i,t}$. Incidentally, the theoretical elasticity of consumption to a transitory shock on net income is neglected because it is approximated as the share of current income in lifetime expected resources, $\phi^{\varepsilon\, BPS}_{\text{net. inc.}} = y_t / ((1+r)a_t + \sum_{s=0}^{T-t} \frac{E_t[y_{t+s}]}{(1+r)^s})$, which is small. The method is more structural than the BPP estimator, since the authors use the theoretical expressions of the elasticities for estimation, and not just the restriction they imply that log-consumption growth is uncorrelated with past shocks.[19]

Bias There are two issues with this method. First, the empirical counterpart of $\phi^{\eta}_{\text{wage}}$ is measured with the BPP estimator of the elasticity to a permanent shock. This estimator is subject to caution as it is not robust to the presence of a correlation with past income shocks, although in the case of permanent shocks I do not quantify the extent of the bias that this correlation causes.

Second, the theoretical expressions $\phi^{\eta\,BPS}_{\text{wage}}$, $\phi^{\eta\,BPS}_{\text{net. inc.}}$ and $\phi^{\varepsilon\,BPS}_{\text{net. inc.}}$ are based on approximations that neglect precautionary behavior. In particular, the expressions of the elasticities to shocks on net income as shares of lifetime expected income and current income in lifetime expected resources are derived from the same approximation of log-consumption growth as used in BPP, and they coincide with the expressions that would hold under perfect foresight, in the absence of uncertainty.[20] The exact expression of the elasticity of consumption to a permanent shock on net income is in fact the share of lifetime expected income minus the response of lifetime expected precautionary consumption growth to a permanent income shock in lifetime expected resources net of lifetime expected precautionary consumption growth. The exact expression of the elasticity of consumption to a transitory shock on net income is in fact the share of current minus the response of lifetime expected precautionary consumption growth to a transitory income shock in lifetime expected resources net of lifetime expected precautionary consumption growth:

$$\phi^{\eta}_{net.inc.} = \frac{\sum_{s=0}^{T-t}\frac{E_t[y_{t+s}]}{(1+r)^s} - \overbrace{(dPG_{i,t}/d\eta_{i,t})}^{\neq 0 \text{ (precaution)}}}{(1+r)a_t + \sum_{s=0}^{T-t}\frac{E_t[y_{t+s}]}{(1+r)^s} - \underbrace{PG_{i,t}}_{>0 \text{ (precaution)}}} \neq \phi^{\eta\,BPS}_{net.inc.}$$

$$\phi^{\varepsilon}_{net.inc.} = \frac{y_t - \overbrace{(dPG_{i,t}/d\varepsilon_{i,t})}^{<0 \text{ (precaution)}}}{(1+r)a_t + \sum_{s=0}^{T-t}\frac{E_t[y_{t+s}]}{(1+r)^s} - \underbrace{PG_{i,t}}_{>0 \text{ (precaution)}}} > \phi^{\varepsilon\,BPS}_{net.inc.},$$

with $PG_{i,t} = \sum_{s=1}^{T-t}\frac{l_{t,s}E_t[\varphi_{t+s-1}]}{(1+r)^s}$ the lifetime expected precautionary consumption growth, and $l_{t,s}$ a deterministic weight. Neglecting precautionary behavior induces an estimation bias in the measure of the Frisch elasticities, as it leads to plugging in expression $\phi^{\eta\,BPS}_{\text{net. inc.}}$ that wrongly

[19]Though BPP note that approximated expression of the elasticities to transitory and permanent shocks on net income in the standard life-cycle model are the shares of current income and lifetime expected income in lifetime expected resources (bottom of p.1897), they do not use these expressions for estimation.

[20]In these two cases, consumption is a constant share of lifetime expected resources. A one unit permanent shock raises lifetime expected resources by an amount equal to lifetime expected income and current income, so its share in lifetime expected resources is the percentage change in consumption caused by a permanent shock when consumption is a constant fraction of lifetime expected resources. Similarly, a one unit transitory shock raises lifetime expected resources by an amount equal to current income.

estimates the sensitivity of consumption to a permanent shock, and expression $\phi^{\varepsilon\, BPS}_{\text{net. inc.}}$ that underestimates its sensitivity to a transitory shock.

Blundell, Pistaferri, and Saporta-Eksten (2016) develop a robustness check in which they allow $\phi^{\eta\, BPS}_{\text{net. inc.}}$, the share of expected lifetime income in expected lifetime resources, to be weighted by a parameter β that the authors estimate and find to be non significant. Yet, such a test would not necessarily capture the bias caused by the difference between $\phi^{\eta\, BPS}_{\text{net. inc.}}$ and its exact expression $\phi^{\eta}_{\text{net. inc.}}$, as their difference is not necessarily a constant fraction of $\phi^{\eta}_{\text{net. inc.}}$. Also, despite the presence of the parameter β, the model still imposes that the elasticities to a transitory shock be zero, and imposes a functional form $f(.)$ on $\phi^{\eta\, BPS}_{\text{wage}}$ that neglects precautionary behavior.

Appendix C Data

The main data source is the PSID, which contains longitudinal information on a representative sample of US households, surveyed every year. It started in 1968 with approximately 3,000 households. Both the original households and their splitoffs have been followed since. The period I consider is 1978-1992.[21] I select out households that are not continuously married over the period, those experiencing a dramatic change in family composition, those headed by a female, those with missing reports on race, education, and region, and those whose head is younger than 30 or older than 65. I also drop some income outliers. The dataset, the period, and the selection are the same as in BPP. The final sample is composed of 15,779 household-year observations from 1,765 households[22]. Among these, there are 12,041 household-year observations for which current log-consumption growth and current log-income growth are simultaneously observed, and 8,958 for which current log-consumption growth, current log-income growth, and log-income growth two periods later are simultaneously observed.

Net income is the taxable family income reported by a household minus its financial income and minus the federal taxes paid on nonfinancial income.[23] Gross income is net income plus taxes. Gross income before transfers is gross income minus transfer income.[24] All three measures are deflated by the contemporaneous Consumer Price Index (CPI).

Nondurable consumption is the sum of annual expenditure on food, alcohol, tobacco, nondurable services, heating fuel, public and private transport (including gasoline), personal care,

[21] The CEX data that is used to impute consumption is difficult to use before 1978. After 1992, a number of the questions used by BPP to build their measure of income are redesigned.

[22] My sample is exactly the same as that of BPP. The number of household-year observations reported differ (they report 17,604 observations) simply because they count observations of log-income, while I report observations of log-income growth, which is the variable used for estimation.

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

[24] Transfer income includes aid to families with dependent children, supplemental security income and other welfare payments, social security income and other retirement, pensions and annuities payments, unemployment benefits, worker's compensations, child support, help from relatives, and other transfer income.

and clothing and footwear, deflated by the CPI. Total consumption is the sum of nondurable consumption plus annual expenditures on durable goods, namely housing (mortgage interest, property tax, rent, other lodging, textiles, furniture, floor coverings, appliances), new and used cars, vehicle finance charges and insurance, car rentals and leases, cash contributions, and personal insurance (life insurance and retirement), deflated by the CPI. Total consumption plus health and education is the sum of total consumption plus annual expenditures on health (insurance, prescription drugs, medical services), and education. As the PSID only reports expenditure on food, these three measures of consumption are imputed from the demographic characteristics of the households and from their food consumption, with the coefficients used for the imputation estimated from the CEX over the same period. Further details are provided in the paper of BPP (section I.B.).

I detrend log-income and log-consumption from the impact of demographic variables by regressing them on dummies for year, year-of-birth, family size, number of children, existence of outside dependent children, education, race, employment status, presence of an additional income recipient that is not the head or his spouse, region, residence in a large city, and interactions between a subset of these demographic characteristics and year and cohort dummies. This follows BPP, except that I add interactions with cohort dummies, which are present in a more recent version of the BPP estimator (the one used by Blundell, Pistaferri, and Saporta-Eksten (2016)).

Appendix D Estimating with a generalized method of moment

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

$$cov(\Delta ln(\tilde{c}_{i,t}),\Delta ln(\tilde{y}_{i,t+2})) = E[\Delta ln(\tilde{c}_{i,t})\Delta ln(\tilde{y}_{i,t+2})] = \phi^{\varepsilon} var(\varepsilon_{i,t}),$$
$$cov(\Delta ln(\tilde{y}_{i,t}),\Delta ln(\tilde{y}_{i,t+2})) = E[\Delta ln(\tilde{y}_{i,t})\Delta ln(\tilde{y}_{i,t+2})] = var(\varepsilon_{i,t}).$$

Thus, the following moment restriction holds:

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

with $X_{i,t} = (\Delta ln(\tilde{c}_{i,t}),\Delta ln(\tilde{y}_{i,t}),\Delta ln(\tilde{y}_{i,t+2}))$ the set of variables involved , and $\phi^{\varepsilon} = E[\phi^{\varepsilon}_{i,t}]$ the parameter involved. This restriction makes it possible to estimate the parameter ϕ^{ε} as the value that minimizes a norm of the sample analog of this moment:

$$\hat{\phi^{\varepsilon}} = \underset{\phi^{\varepsilon}}{\operatorname{argmin}} \left(\frac{1}{N} \sum_{n=1}^{N} g(X_n, \phi^{\varepsilon}) \right)^{\mathsf{T}} \hat{W} \left(\frac{1}{N} \sum_{n=1}^{N} g(X_n, \phi^{\varepsilon}) \right),$$

with N the number of household-year observations (i,t) at which the three variables $\Delta ln(\tilde{c}_{i,t})$, $\Delta ln(\tilde{y}_{i,t})$, and $\Delta ln(\tilde{y}_{i,t+2})$ are observed, and $\hat{W}$ a weighting matrix. The matrix is chosen so the estimation of the standard error is robust to arbitrary within-household correlations and robust to heteroskedasticity, that is, to the fact that the residuals $g(X_{i,t}, \phi^{\varepsilon}) - E[g(X_{i,t}, \phi^{\varepsilon})]$ are not drawn from the same distribution.

Appendix E Robustness

E.1 Anticipation of the shocks

I generalize the statistical model by allowing part of the realizations of the permanent and transitory shocks at t to be anticipated at previous periods $t-s$ and $t-k$:

$$\eta_{i,t} = \eta_{i,t}^{surp} + \eta_{i,t-s}^{ant,t},$$
$$\varepsilon_{i,t} = \varepsilon_{i,t}^{surp} + \varepsilon_{i,t-k}^{ant,t}.$$

Each type of shock writes as the sum of a surprise component and an anticipated component whose value realizes at t but is known at $t-s$ or $t-k$. The anticipated component is uncorrelated with the current and past surprise components and with the past anticipated components.

Earliest anticipation period From the moments presented in Table 1, log-consumption growth at t no longer covaries with future log-income growth at $t+3$ or later, which is informative about how early each type of shock can be anticipated. Indeed, in the presence of anticipation, the covariance between log-consumption growth and future log-income growth at $t+3$ is:

$$\begin{aligned} 0 = cov(\Delta ln(\tilde{c}_{i,t}), \Delta ln(\tilde{y}_{i,t+3})) = & \underbrace{cov(\Delta ln(\tilde{c}_{i,t}), \eta_{i,t+3-s}^{ant,t+3})}_{\neq 0 \text{ if } s>2} + \underbrace{cov(\Delta ln(\tilde{c}_{i,t}), \varepsilon_{i,t+3-k}^{ant,t+3})}_{\neq 0 \text{ if } k>2} \\ & - (1-\theta) \underbrace{cov(\Delta ln(\tilde{c}_{i,t}), \varepsilon_{i,t+2-k}^{ant,t+2})}_{\neq 0 \text{ if } k>1} - \theta \underbrace{cov(\Delta ln(\tilde{c}_{i,t}), \varepsilon_{i,t+1-k}^{ant,t+1})}_{\neq 0 \text{ if } k>0}. \end{aligned}$$

Apart from a knife-edge case in which the effects would perfectly compensate each other, the earliest period at which a permanent shock can be anticipated that is consistent with the covariance above being zero is $s=2$. The earliest period at which a transitory shock can be anticipated is $k=0$, that is, a transitory shock cannot be anticipated beforehand.[25]

[25] If households are myopic or constrained, it is possible that they anticipate shocks earlier but simply do not

Bias from anticipation Applying the robust estimator to a model in which permanent shocks are partly anticipated gives:

$$\hat{\phi}^{\varepsilon} = \frac{cov(\Delta ln(\tilde{c}_{i,t}), -\Delta ln(\tilde{y}_{i,t+2}))}{cov(\Delta ln(\tilde{y}_{i,t}), -\Delta ln(\tilde{y}_{i,t+2}))} = \phi^{\varepsilon} - \underbrace{\frac{cov(\Delta ln(\tilde{c}_{i,t}), \eta^{ant,t+2}_{i,t+2-s})}{\theta var(\varepsilon_{i,t})}}_{=0 \text{ if } s=1 / /> 0 \text{ if } s=2} \leq \phi^{\varepsilon}.$$

The denominator of the estimator is unaffected because whether a household can anticipate future income shocks or not does not affect its current log-income growth, which only correlates with future log-income growth through the realization of the transitory shock. The numerator can be affected, however, because information about its future income can impact the current consumption decision of a household, inducing a correlation between current log-consumption growth and future log-income growth through the anticipated component of future log-consumption growth. If the anticipated component is known only one period in advance, though, the numerator remains unaffected: no component of the instrument is anticipated yet at t, so the current transitory shock is still the only variable through which current log-consumption growth and the instrument covary. If the anticipated component is known two periods in advance, and if log-consumption is positively correlated with current information about the realization of future permanent income, then the robust estimator is a conservative measure of the elasticity and underestimates it.

E.2 Serial correlation in measurement error

I consider a more general model in which measurement error can be correlated over time. Measurement error ζ^y is no longer orthogonal to its past values but such that: $\zeta^y_{i,t} = \dot{\zeta}^y_{i,t} + \nu \dot{\zeta}^y_{i,t-1}$, with ν a parameter measuring the strength of the serial correlation. Thus, log-income growth depends on past measurement error $\dot{\zeta}^y$ up to two periods ago:

$$\Delta ln(\tilde{y}_{i,t}) = \eta_{i,t} + \varepsilon_{i,t} - (1-\theta)\varepsilon_{i,t-1} - \theta\varepsilon_{i,t-2} + \dot{\zeta}^y_{i,t} - (1-\nu)\dot{\zeta}^y_{i,t-1} - \nu\dot{\zeta}^y_{i,t-2}$$

The robust estimator is:

$$\hat{\phi}^{\varepsilon} = \frac{cov(\Delta ln(\tilde{c}_{i,t}), -\Delta ln(\tilde{y}_{i,t+2}))}{cov(\Delta ln(\tilde{y}_{i,t}), -\Delta ln(\tilde{y}_{i,t+2}))} = \phi^{\varepsilon} \times \underbrace{\frac{\theta var(\varepsilon_{i,t})}{\theta var(\varepsilon_{i,t}) + \nu var(\dot{\zeta}^y_{i,t})}}_{<1} < \phi^{\varepsilon}.$$

respond to them. In that case, the fact that the covariance is zero does not imply that shocks are not anticipated earlier than $s=2$ and $k=0$. Yet, in that case, an early anticipation of the shocks does not bias the estimation since consumption does not respond to the anticipated component of a shock before it actually realizes.

The presence of serial correlation in measurement error leads to an overestimation of the variance of the current transitory shock, because the variance of measurement error is mistaken as variance of the transitory shock: the denominator measures $\theta(var(\varepsilon_{i,t}) + (\nu/\theta)var(\dot{\zeta}_{i,t}))$ instead of $\theta var(\varepsilon_{i,t})$. The estimation of the covariance between log-consumption growth and a current transitory shock is unaffected because consumption does not respond to measurement error. Thus, in the presence of serial correlation, the robust estimator is conservative and underestimates the elasticity of consumption to a transitory shock because it overstates the importance of the shocks. If I additionally allow measurement error in consumption to correlate with measurement error in income, the robust estimator would overestimate both the covariance between log-consumption growth and a transitory shock and the variance of the transitory shock, still underestimating the elasticity to a transitory shock as long as the elasticity of consumption to measurement error in income is smaller than its elasticity to a transitory shock and that ν is small when compared to θ.[26]

E.3 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 means that a permanent shock η_t still affects the value of permanent income at each period in the future until the rest of the household's lifetime, but the effect of this permanent shock now depreciates at a rate $(1-\rho)$ instead of affecting all values of future permanent income in a the same way. As a result, the log-income growth of a household at t depends, not only on the past transitory shocks it has received at $t-1$ and $t-2$, but also on all the permanent shocks it has received in the past:

$$\begin{aligned}\Delta ln(\tilde{y}_t) =& \eta_t - (1-\rho)\eta_{t-1} - (1-\rho)\rho\eta_{t-2} - ... - (1-\rho)\rho^{t-2}\eta_1 - (1-\rho)\rho^{t-1}p_0 \\ & + \varepsilon_t - (1-\theta)\varepsilon_{t-1} - \theta\varepsilon_{t-2}.\end{aligned}$$

The instrument, future log-income growth at $t+2$, no longer identifies only the effect of the current transitory shock because it covaries with current log-consumption growth through the current transitory shock but also through the current and past permanent shocks. Since the correlation between log-consumption growth and past permanent shocks is undetermined, it is difficult to predict the direction of the bias caused by this depreciation. Yet, Kaplan and Vi-

[26]In that case, the robust estimator measures $\hat{\phi}^{\varepsilon} = \frac{\theta cov(\Delta ln(c_{i,t}),\varepsilon_{i,t}) + \nu cov(\Delta ln(c_{i,t}),\dot{\zeta}^y_{i,t})}{\theta var(\varepsilon_{i,t}) + \nu var(\dot{\zeta}^y_{i,t})} < \frac{cov(\Delta ln(c_{i,t}),\varepsilon_{i,t})}{var(\varepsilon_{i,t})} = \phi^{\varepsilon}$ if $\frac{cov(\Delta ln(c_{i,t}),\dot{\zeta}^y_{i,t})}{var(\dot{\zeta}^y_{i,t})} < \frac{1-\theta}{\nu}\frac{cov(\Delta ln(c_{i,t}),\varepsilon_{i,t})}{var(\varepsilon_{i,t})}$.

olante (2010) show that if the value of ρ is known, it is possible to obtain a consistent estimator by substituting log-income growth with its quasi-difference $\Delta^{\rho} ln(\tilde{y}_{i,t}) = ln(\tilde{y}_t) - \rho ln(\tilde{y}_{t-1})$ in the estimating moments. The estimator is:

$$\hat{\phi}^{\varepsilon,\rho} = \frac{cov(\Delta ln(\tilde{c}_{i,t}), \Delta^{\rho} ln(\tilde{y}_{i,t+2}))}{cov(\Delta^{\rho} ln(\tilde{y}_{i,t}), \Delta^{\rho} ln(\tilde{y}_{i,t+2}))} = \phi^{\varepsilon}.$$

Table 5: Elasticity $\phi^{\varepsilon,\rho}$ and MPC lower bound $\underline{MPC}^{\varepsilon,\rho}$ when permanent shocks depreciate

Depreciation	$\rho = 1$	$\rho = 0.97$	$\rho = 0.94$	$\rho = 0.91$
$\phi^{\varepsilon-\rho}$	0.622	0.626	0.648	0.690
	(0.249)	(0.258)	(0.276)	(0.306)
$\underline{MPC}^{\varepsilon-\rho\ total}$	0.450	0.458	0.483	0.535
	(0.229)	(0.241)	(0.265)	(0.309)
Obs.	6,337	6,337	6,337	6,337

Note: See main text

Table 4 presents the results obtained with such an estimator, for different values of ρ. The first column corresponds to the baseline case in which $\rho = 1$. The point estimate of the elasticity of consumption then increases as ρ decreases below one: it moves from 0.54 to 0.55, 0.58, and 0.64 when ρ moves from 1 to 0.97, 0.94, and 0.91. It means that, if in fact permanent shocks depreciate over time, the robust estimator is conservative and underestimates the elasticity.

Appendix F Alternative specifications

F.1 Demographic characteristics

Table 6: Elasticity ϕ^{ε} - Alternative sets of detrending variables

Detrending	Baseline	Baseline + past char. at $t-2$	Only empl. + past char. at $t-1$	BPP set + past char. at $t-1$	BPP set
ϕ^{ε}	0.597	0.657	0.726	0.623	0.425
	(0.246)	(0.348)	(0.356)	(0.257)	(0.299)
Obs.	7,600	6,349	7,600	7,600	8,958

Note: See main text.

Table 7: Elasticity ϕ^{ε} - Alternative clusters

Cluster level	Household	Coh. $\times$ edu.	Year $\times$ edu.
ϕ^{ε}	0.597 (0.246)	0.597 (0.231)	0.597 (0.209)
Num. of clusters	1,454	116	24
Obs.	7,600	7,600	7,600

Note: Standard errors in parentheses are adjusted for heteroskedasticity. The first line of this table reports GMM estimates of the average elasticity of nondurable consumption to a transitory shock on net income and estimates of their standard errors allowing for arbitrary correlation within a household, within households with the same level of education and the same year-of-birth, and within households with the same level of education observed on the same year. The second line reports the number of clusters within which arbitrary correlation is allowed. The third line reports the number of household-years for which the estimating moment is observed.

Table 8: Elasticity ϕ^{ε} - Alternative measures of consumption

	Nondur.	Food	Total excl. edu. & health	Total
ϕ^{ε}	0.597 (0.246)	0.335 (0.205)	0.717 (0.300)	0.714 (0.296)
Obs.	7,600	7,613	7,600	7,600
Add. dem.	yes	yes	yes	no
Moments	(1.6)	(1.6), (1.10)	$(1.6)_t$, $(1.10)_t$ $\forall t$ + others	$(1.6)_t$, $(1.10)_t$ $\forall t$ + others

Note: Standard errors in parentheses are adjusted for arbitrary within-household correlations and heteroskedasticity. The number of observations is the number of household-year pairs for which either (1.6) or (1.10) are observed (in the second column I impose that both are observed so the sample is exactly the same as in the first column).

Table 9: Elasticity ϕ^{ε} - Alternative measures of income

	Net income (earn.+trans.-tax.)	Gross income (earn.+trans.)	Earnings (male + fem. earn.)
ϕ^{ε}	0.597 (0.246)	0.537 (0.216)	0.301 (0.146)
Obs.	7,600	7,600	7,224
Add. dem.	yes	yes	yes
Moments	(1.6)	(1.6)	(1.6)

Note: Standard errors in parentheses are adjusted for arbitrary within-household correlations and heteroskedasticity. The number of observations is the number of household-year pairs for which either (1.6) or (1.10) are observed (in the second column I impose that both are observed so the sample is exactly the same as in the first column).