Durable Goods, Borrowing Constraints and Consumption Insurance

Enzo A. Cerletti
CEMFI

Josep Pijoan-Mas*
CEMFI and CEPR

November 15, 2011

Abstract

In this paper we study the transmission of income shocks into nondurable consumption in the presence of durable goods. We build a two-good model to characterize the interaction of durability of goods, durability of shocks, and borrowing constraints as determinants of shock transmission. We show that borrowing constraints lead to a basket rebalancing effect, which biases the measures of insurance based on the consumption responses to shocks. The sign of this rebalancing depends critically on the persistence of the shock. We calibrate the model economy to the US in order to measure the size of this bias.

JEL classification: E21, D91, D12
Keywords: Consumption Insurance, Durable Goods, Incomplete Markets, Borrowing Constraints, Persistence of Income Shocks

*Postal address: CEMFI, Casado del Alisal 5, 28014 Madrid, Spain. E-mail: cerletti@cemfi.edu.es, pijoan@cemfi.es

1 Introduction

The standard life-cycle model of consumption predicts that, if households had access to perfect insurance markets against income shocks, they would choose to smooth out completely consumption fluctuations. As a result, at least since Hall and Mishkin (1982), the response of expenditures in non-durable consumption goods to income shocks has been used to measure the amount of insurance available to private households. For instance, the recent work by Blundell, Pistaferri, and Preston (2008) finds that consumption expenditures hardly react to transitory shocks, whereas 64 percent of permanent income shocks translate into consumption. These facts suggest that transitory shocks are very easy to insure against but permanent shocks are not. According to Kaplan and Violante (2010), a standard life-cycle model of consumption —with borrowing and saving as the only source of insurance— would rightly predict the very small response of consumption expenditures to transitory income shocks. But it would overpredict the large empirical response of consumption to permanent shocks, suggesting that households must have access to forms of insurance beyond self-insurance. This result is in line with the *excess smoothness* of consumption data to (permanent) income changes, which has been largely discussed in the consumption literature.[1]

The use of consumption responses to income shocks as a measure of insurance ignores substitution with durable goods. This is not a problem under the assumptions of (a) homothetic preferences, (b) absence of adjustment costs in durable goods, and (c) absence of binding borrowing constraints. In this situation, a standard consumption model predicts that the optimal composition of the basket between durable and non-durable goods is constant and does not change with income shocks. Hence, the response of non-durable expenditures to income shocks is a sufficient statistic for the changes in the marginal utility of consumption and hence it does reflect the level of insurance. However, as first shown by Chah, Ramey, and Starr (1995), binding borrowing constraints change this. Households that hit the borrowing constraint shift their consumption basket away from durable goods, as their durability is valued less when households would like to bring consumption from the future to the present and cannot do so. A positive income innovation that is not too persistent diminishes the severity of the borrowing constraint, and hence spurs a rebalancing from non-durable goods towards durable goods. This fact is used by Luengo-Prado (2006) to show that the excess smoothness puzzle may be partly explained by the rebalancing of durable and non-durable goods.

In this paper, we study the response to income shocks of consumers that care for both

[1] For an early formulation of the excess smoothness puzzle see Campbell and Deaton (1989).

durable and non-durable goods. We keep homothetic preferences over the two types of consumption goods, we abstract from adjustment costs in the stock of durable goods, and we focus on the borrowing constraints. Given these assumptions, we think of durable goods as cars, furniture, home appliances, and the like, but we exclude housing in our main exercise. We characterize theoretically the bias in the consumption insurance measures due to rebalancing of durable and non-durable goods. We show that the size of the bias depends negatively on the persistence of the income shocks, positively on the durability of the goods, and negatively on the fraction of the durable goods that can be collateralized. We also show that with income shocks that are very persistent (permanent) the sign of the bias can be reversed. This does not happen if a household has an optimal consumption profile that increases over the life-cycle, as then a permanent positive shock to labor earnings also alleviates the borrowing constraint. However, if a household has a desired consumption profile that falls over time, a permanent income shock makes borrowing constraints more severe and then the sign of the bias is reversed. In this situation, true insurance of consumption to permanent shocks is larger than what is typically measured in the data.

In order to measure which fraction of the transmission of income shocks to consumption expenditures is due to lack of insurance and which fraction is due to rebalancing between durable and non-durable goods, we calibrate our economy to household level data on consumption expenditures and wealth holdings. This allows us to recover key quantitative parameters as the down payment against the basket of durable goods, the severity of borrowing constraints and the shape of the desired life-cycle consumption profiles. Then, we check the model ability to reproduce the observed response of expenditure in non-durable composition to income shocks at different ages, and quantify the bias in conventional measures of consumption insurance.

In our main calibration we find that, on aggregate, the average transmission of income into non-durable consumption is a good measure of the degree of insurance. The response of non-durable consumption to unexpected income changes is 13.1 percent when they are transitory and 74.3 percent when they are permanent. A very small part of these responses is due to the rebalancing between durable and non-durable goods. In particular, the stock of durable goods responds 2 percentage points more to transitory shocks and 3.3 percentage points less to permanent shocks. Hence, transitory shocks generate a rebalancing towards durable goods and permanent shocks generate a rebalancing away from durable goods. For households with a head below 40 years of age the transmission of income shocks to consumption is larger and includes more rebalancing. In particular, for this age group, the model predicts that 20.6 percent of transitory shocks get transmitted into non-durable consumption and the transmission into durable goods is 5 percentage points higher, while for permanent shocks the transmission is 98 for non-durables and 8 percentage points less for durable goods.

Finally, we also compare the transmission of shocks into consumption of our model with durables to a model without durables. We calibrate both models to the same amount of total household wealth and the same fraction of households hitting the borrowing constraints. The two models differ in the possible substitution of goods upon income shocks, but also in the higher ability to borrow in the model with durables (due to the collateral value of the durable goods) and in the timing of the wealth accumulation, which is earlier in life in the model with durables. We find that these two additional differences revert the previous results and hence, the transmission of income shocks into non-durable consumption understate the true level of insurance against transitory shocks and overstate the insurance against permanent shocks.

It would be interesting to test empirically whether unexpected income changes drive responses in the ratio of durable and non-durable goods which are different for constrained and unconstrained households. Testing the rebalancing effect in the data, however, is notoriously difficult. Several authors find that expenditure in durables increase (fall) more than expenditures in non-durables upon the arrival of positive (negative) shocks. For instance, Browning and Crossley (2009) show that among a sample of Canadian unemployed workers, those with lower unemployment benefits reduce expenditure on durable goods more, and do so more for those goods with higher durability. Johnson, McClelland, and Parker (2010) look at the consumer responses to the reception of the checks of the Economic Stimulus Payments of 2008 in the US. They also find the response to be larger for durable goods than for non-durable goods. Aaronson, Agarwal, and French (2008) find that households affected by a minimum wage hike increase expenditure in durables much more than in non-durables, while increasing collateralized debt at the same time. However, this evidence is not easy to interpret. As pointed out by Bils and Klenow (1998), to achieve a given increase in the stock of durables one needs to increase expenditure more when the durability of the good is higher (the depreciation rate lower). Hence, evidence of durable expenditure reacting more to transitory income shocks than non-durable expenditure does not need to reflect a rebalancing of durable and non-durable goods.

The remaining of the paper is organized as follows. We describe the basic model in Section 2 and we calibrate it in Section 3. Then, in Section 4 we discuss how to measure the bias in the conventional measure of insurance to shocks and present our main quantitative findings. In Section 5 we discuss several alternative calibrations. Finally, Section 6 concludes.

2 The Model

2.1 Timing

Households live for an uncertain life span of at most T periods. The probability of surviving between age t and $t+1$ is given by $\pi_{t,t+1}$, with $\pi_{T,T+1} = 0$. We define $\pi_t = \prod_{j=1}^{t-1} \pi_{j,j+1}$ as the unconditional probability of a newborn surviving up to age t. Their lives are divided in two main periods: working life and retirement. Households are born as working adults, so working life lasts from age $t = 1$ to age $T = T_R$, at which retirement is mandatory. Hence, T_R is the unique retirement age in this economy. From $t = T_R + 1$ to $t = T$, they live as retirees.

2.2 Preferences and decisions

Households have time-separable preferences defined over streams of consumption of non-durable goods C_t and service flows of durable goods D_t, which are assumed to be proportional to its stock. Time preferences are captured by the discount factor β, so that at any age t households discount the utility of age $t+j$ at a rate $\beta^j \frac{\pi_{t+j}}{\pi_t}$. Under these assumptions, the lifetime objective function of a household is given by

$$E_0 \sum_{t=1}^{T} \beta^{t-1} \pi_t U\left(C_t, D_t\right) \tag{1}$$

We assume CRRA preferences over a CES aggregator of non-durable consumption and services from durable goods:

$$U\left(C_t, D_t\right) = \frac{\left[\gamma C_t^{\epsilon} + \left(1-\gamma\right) D_t^{\epsilon}\right]^{\frac{1-\sigma}{\epsilon}}}{1-\sigma}$$

where $\sigma > 0$ measures the degree of risk aversion, $\eta < 1$ measures the elasticity of substitution between goods and $1 > \gamma > 0$ captures the weight of each type of consumption in households' preferences.

Households seek to maximize their lifetime utility by choosing each period t how much to spend in non-durable consumption, C_t, how many durable goods to buy (or sell), I_t, and how many financial assets to save (or borrow) for the next period, A_{t+1}, subject to a set of constraints described below.

2.3 Durable goods

The stock of durable goods evolve according to the following law of motion:

$$D_t = (1-\delta)\, D_{t-1} + I_t \tag{2}$$

where $(1-\delta)$ is the fraction of the end-of-period stock at $t-1$ that remains providing utility at t. Notice that the utility function in (1) depends on the end-of-period stock of durables, D_t, after period t purchases and sales, I_t. Durable goods have a relative price of P, with non-durable goods acting as the numeraire.

2.4 Financial assets and borrowing constraints

Households save and borrow through a perfectly competitive annuity market. We denote A_t the amount saved or borrowed in annuities, which yields a constant net return r, and r_t is the age-specific interest rate corrected by survival probabilities, that is, $r_t = (1+r)/\pi_{t,t+1} - 1$. We model borrowing constraints by restricting a measure of end-of-period households' net worth to be above a threshold $\underline{A}_t$. Moreover, the measure of net worth only incorporates the value of the end-of-period stock of durables up to a fraction $0 \leq \theta \leq 1$, implying a limited role of durables as collateral. Hence, financial assets are bounded below by,

$$(1+r_{t+1})A_{t+1} + \theta\,(1-\delta)\, PD_t \geq \underline{A}_t \tag{3}$$

We will consider a specification for $\underline{A}_t$ reflecting financial market imperfections and heterogeneous borrowing capacity at the same time, by linking $\underline{A}_t$ to the natural borrowing limit. At age t, the lower bound to the present discounted value of labour income is given by

$$\underline{Y}_t + \sum_{j=t+1}^{T}\left(\prod_{k=t+1}^{j}\frac{1}{1+r_k}\right)\underline{Y}_j \tag{4}$$

where $\underline{Y}_j$ denotes the lowest possible income for age j given the information available at age t. Therefore, expression (4) is the maximum amount that can be repaid at age t with probability one, the natural borrowing limit for an agent in the model economy. We can use (4) to define the borrowing constraint as

$$\underline{A}_t = -\alpha\left[\underline{Y}_{t+1} + \sum_{j=t+2}^{T}\left(\prod_{k=t+2}^{j}\frac{1}{1+r_k}\right)\underline{Y}_j\right] \tag{5}$$

where $1 \geq \alpha \geq 0$ is the parameter determining the overall tightness of the borrowing constraints. Thus, expression (3) means that the agent's relevant net worth at the beginning of age $t+1$, which is decided at age t, must be equal to or greater than a fraction α of the present value of all the resources it can generate from age $t+1$ to T with certainty.[2]

An extreme case is $\alpha = 0$ and $\theta = 0$, which precludes borrowing altogether. Another interesting case to consider is $\alpha = 0$ and $\theta > 0$, a case in which collateralized debt is the only form of borrowing allowed. The particular case of $\alpha = 0$ and $\theta = 1$ can be rationalized as emerging from a limited commitment setup, in which the penalty for defaulting is the seizure of the whole stock of durables, as in Fernandez-Villaverde and Krueger (2010).

2.5 Labor income and pension income

Income Y_t while working is given by a stochastic process to be defined below. After age T_R, households receive an age-invariant payment Y_{T_R} from social security. This payment is household-specific and it is made a function of its entire labour income history. At any age, we summarize past earnings history or social security wealth in the variable $H_t = \frac{1}{t}\sum_{j=1}^{t} Y_j$, or recursively:

$$H_t = \frac{(t-1)H_{t-1} + Y_t}{t} \tag{6}$$

Thus, pension payments can be written as a function of H_{T_R}. Making pension payments a function of the history of past income is important. As it will be shown in later sections, the amount of rebalancing between durable and non-durable goods depends on the persistence of the income shocks. And the persistence of the income shocks depends on how much of them is translated into pension income, as well as on the exact nature of the stochastic process.

We assume that the stochastic process governing the log of wages, y_t, can be represented as the sum of a random walk, z_t, a purely transitory shock, ε_t, and a deterministic,

[2]The agent can die before T. However, the discount factor is the corrected interest rate, which incorporates the survival probabilities. Hence, a financial intermediary lending to the continuum of households would always recover the amount lent plus the riskless rate r.

age-specific mean, μ_t:

$$\begin{aligned} y_t &= \mu_t + z_t + \varepsilon_t \\ z_t &= z_{t-1} + \eta_t \\ \varepsilon_t &\sim \mathcal{N}(0, \sigma_\varepsilon) \\ \eta_t &\sim \mathcal{N}(0, \sigma_\eta) \\ z_0 &\sim \mathcal{N}(0, \sigma_{z_0}) \end{aligned} \tag{7}$$

where η_t is a permanent shock to labour income.

2.6 Period budget constraint

With all the elements defined we can construct the budget constraint during working life as:

$$C_t + PI_t + A_{t+1} \leq (1 + r_t)A_t + Y_t \tag{8}$$

and during retirement:

$$C_t + PI_t + A_{t+1} \leq (1 + r_t)A_t + Y^R(H_{T_R}) \tag{9}$$

2.7 Choices

Summing up, households choose the sequences $\{C_t\}_{t=1}^T$, $\{I_t\}_{t=1}^T$, and $\{A_t\}_{t=2}^T$ to maximize (1), subject to a sequence of T_R constraints of the form in (8) and $T - T_R$ constraints of the form in (9), the laws of motion defined by (2) and (6), the borrowing constraints (3), the stochastic process for labor income defined in (7), and some initial conditions A_1, D_0, and z_0.[3]

2.8 State variable transformation

It may be useful to rewrite the period budget constraint in terms of total resources available to the household at a given time, including the amount that it will be able to borrow. Let's define cash on hand at the beginning of period t as:

$$\tilde{X}_t = (1 + r_t)A_t + (1 - \delta)\, PD_{t-1} + Y_t - \frac{\underline{A}_t}{(1 + r_{t+1})} \tag{10}$$

[3] From its definition, it is understood that $H_0 = 0$.

and

$$\tilde{A}_{t+1} = A_{t+1} - \frac{\underline{A}_t}{(1 + r_{t+1})} \tag{11}$$

We can then write the period budget constraint as

$$C_t + PD_t + \tilde{A}_{t+1} \leq \tilde{X}_t \tag{12}$$

Then, we can rewrite (3) as

$$\tilde{A}_{t+1} + \frac{\theta (1 - \delta) PD_t}{(1 + r_{t+1})} \geq 0 \tag{13}$$

where a negative $\tilde{A}_{t+1}$ can be interpreted as collateralized debt.

2.9 Lagrangian formulation

The problem can be summarized in the following Lagrangian:

$$\begin{aligned} L &= E_0 \sum_{t=1}^{T} \beta^t \pi_t U(C_t, D_t) \qquad (14) \\ &+ E_0 \sum_{t=1}^{T} \beta^t \pi_t \mu_t \left[\tilde{X}_t - PD_t - C_t - \tilde{A}_{t+1} \right] \\ &+ E_0 \sum_{t=1}^{T} \beta^t \pi_t \lambda_t \left[\tilde{A}_{t+1} + \frac{\theta (1 - \delta) PD_t}{1 + r_{t+1}} \right] \end{aligned}$$

where the multiplier λ_t reflects the value of marginally relaxing the borrowing constraint, while μ_t is the standard multiplier associated to the period budget constraint. The first order conditions for an optimum are given by

$$U_C(C_t, D_t) = \mu_t \tag{15}$$

$$\beta \pi_{t,t+1} (1 + r_{t+1}) E_t [\mu_{t+1}] = \mu_t - \lambda_t \tag{16}$$

$$\frac{1}{P} U_D(C_t, D_t) + \beta \pi_{t,t+1} (1 - \delta) E_t [\mu_{t+1}] = \mu_t - \lambda_t \frac{\theta (1 - \delta)}{1 + r_{t+1}} \tag{17}$$

Equation (15) equalizes the shadow value of resources within the period to the marginal utility of consumption. Equation (16) is the standard Euler equation that describes the law of motion of the shadow value of wealth. It states that the value of investing one unit of non durable consumption in the financial asset (left hand side) must equal its cost (right hand side). This cost is the shadow value of resources today minus the value of relaxing the borrowing constraint, which is given by the multiplier λ_t. Hence, when the borrowing constraint binds ($\lambda_t > 0$) the expected growth of marginal utilities is lower and the expected consumption growth is higher. Equation (17) drives the choice of durable goods. The left hand side states the value of buying today one unit of the durable good, which is the utility flow of the durable good today plus the value tomorrow of the undepreciated stock. The right hand side is the cost, which is given by the shadow value of the resources used to buy the unit of durable good minus the value of relaxing the borrowing constraint. The value of the durable good in relaxing the borrowing constraint is given by the multiplier λ_t times the fraction of the durable good that can be collateralized, and hence that it is useful in relaxing the borrowing constraint.

2.9.1 Optimal basket of durable and non-durable goods

Combining the optimality equations we obtain:

$$\frac{U_D\left(C_t, D_t\right)}{U_C\left(C_t, D_t\right)} = \left(\frac{r_{t+1} + \delta + \frac{\lambda_t}{\mu_t}\left(1-\delta\right)\left(1-\theta\right)}{1 + r_{t+1}}\right) P \tag{18}$$

Whenever the borrowing constraints do not bind at t, we have that $\lambda_t = 0$ and this expression reduces to the standard condition

$$\frac{U_D\left(C_t, D_t\right)}{U_C\left(C_t, D_t\right)} = \left(\frac{r_{t+1} + \delta}{1 + r_{t+1}}\right) P \tag{19}$$

This states that the marginal rate of substitution between durable and non-durables is equal to their relative price times the user cost of durables. Hence, the ratio between marginal utilities is independent of individual level variables and will be equalized across households of a given age t. Note that, while an income shock can translate into the growth rate of consumption, it can not have a differentiated impact on each type of goods. The intuition is that, without any restriction to adjust the C_t/D_t ratio, and given the isoelastic nature of the utility function, only the level of consumption bundle reacts to shocks, but not its composition. Thus, including durable goods in this simple manner has no consequences for the study of consumption responses to income shocks. In fact, under the assumed utility function the consumption ratio is also kept constant at

$$\left(\frac{C_t}{D_t}\right)^{1-\epsilon} = \frac{\gamma}{1-\gamma}\left(\frac{r_{t+1} + \delta}{1 + r_{t+1}}\right) P \tag{20}$$

and this equation can be used in the Euler equation to derive the same expression for non-durable consumption growth as in previous studies that omit durable goods, such as Blundell, Pistaferri, and Preston (2008).

In the case of binding borrowing constraints, this result no longer holds. With $\lambda_t > 0$ the user cost of durables is larger than with $\lambda_t = 0$ and so are the marginal rate of substitution and the ratio C_t/D_t. The new term captures the opportunity cost of the durable good in the future —when consumption has less value than in the present— minus the value of the durable good as collateral. When the borrowing constraint is binding, the value of the $(1-\delta)$ units of the stock of durable good that are left tomorrow falls because the household would like to bring consumption from the future to the present. Hence, it is less worthy to buy a durable good today and the ratio C_t/D_t goes up. However, this effect is partly offset by the collateral services of the durable good, which depend on the fraction $(1-\delta)\,\theta$ that can be collateralized. The more severe the value of the borrowing constraint (higher λ_t/μ_t) and the smaller the value of the durable good as collateral (lower θ), the higher the ratio C_t/D_t. In the limit, if the residual value of the durable good expenditure can be collateralized completely, $\theta = 1$, then the optimal ratio C_t/D_t is as in the case without binding borrowing constraints. In the case that the durable good could be used for non-collateralized loans ($\theta > 1$), then we would have that with binding borrowing constraints the share of durable goods would be larger than in the case without borrowing constraints.

Now, how does the basket of consumption goods change with income shocks? This will depend on how the income shocks affect the severity of the borrowing constraints and hence the ratio of multipliers $\frac{\lambda_t}{\mu_t}$, which can be written in terms of the expected path for consumption:

$$\frac{\lambda_t}{\mu_t} = 1 - E_t\left[\frac{\mu_{t+1}}{\mu_t}\right]\beta(1+r_{t+1})$$

A purely transitory shock alleviates the borrowing constraint: it increases current income without changing expected future income. Hence, the household raises current consumption, which decreases the current marginal utility of consumption μ_t without affecting the expected marginal utility of consumption tomorrow $E_t\mu_{t+1}$. Then λ_t/μ_t falls and there is a rebalancing towards durable goods and away from non-durable goods. The size of the rebalancing will depend on the product $\beta(1+r_{t+1})$, on the durability $(1-\delta)$ of the good, and on the share of the durable good that can be collateralized θ. Hence, we need to measure these parameters in order to provide a quantitative answer on the importance of the rebalancing between durable and non durable goods. When an income shock has some persistence, these effects are mitigated. An income shock with some persistence increases current income, but it also increases expected future income. This changes $E_t\left[\mu_{t+1}/\mu_t\right]$

by a lesser extent and hence makes the rebalancing smaller.

With a permanent shock, both current income and future income increase in the same manner. A household with a desired flat profile of consumption would not see and change in the severity of its borrowing constraints. If a household desires an increasing consumption profile, then the permanent income shock still alleviates the borrowing constraints. However, when the household wants to borrow from the future to bring consumption into the present, the increase in expected future income means that his desired consumption today increases more than his current income and hence the severity of the borrowing constraints increases. In this situation, a permanent income shock generates an increase in the ratio λ_t/μ_t and hence there would be a rebalancing towards non-durable goods.

3 Calibration

Our quantitative exercise is the following. We want to choose values for the key parameters α, δ, γ, and β such that the model is consistent with the observed life-cycle profiles of assets, expenditures on durable goods and expenditures on non-durable goods. The remaining parameters are measured directly from the data or fixed to standard values. Then, we will ask the model its predictions for the life-cycle profiles of the transmission coefficients of shocks to expenditures and we will assess quantitatively how much of the transmission is due to rebalancing between durable and non-durable goods and how much is due to lack of insurance.

We use three main data sources at the household level: the Panel Study of Income Dynamics (PSID), the Consumer Expenditure Survey (CEX), and the Survey of Consumer Finance (SCF). Our PSID sample corresponds to the one contained in the Blundell, Pistaferri, and Preston (2008) CEX-PSID imputed dataset. As for the CEX, we work with two datasets: the quarterly panel used in Heathcote, Perri, and Violante (2010), and the series of annual cross-sections described in Harris and Sabelhaus (2000). Unless stated otherwise, the reference period is 1980-1992, the time interval covered by Blundell, Pistaferri, and Preston (2008).

We classify the different expenditure categories in the CEX as either durable or non-durable. Durable goods include cloth, jewelry, furniture, household appliances, vehicles and spare parts, books, and sport and recreational equipment, but exclude housing. Non-durables include food and other household supplies, household utilities, services, public transport fees, fuel and tolls and education expenditures.[4] We exclude health expendi-

[4] Cloth is considered a semi-durable, and has often been included among nondurables in previous studies. Treating it as nondurable has no effect in our quantitative exercise.

tures from the analysis.[5] We also use aggregate data on durables from the 2011 revision of *Fixed Reproducible Tangible Wealth* by the Bureau of Economic Analysis. Our definition of durables in the aggregate data closely follows the one in the micro data, and it's basically obtained from subtracting therapeutical equipment from the total stock of consumer durables.

3.1 Timing and demographics

A period is a year. We assume households are born to working life at age 25 and retire at age 60. Certain death takes place at age 95. This implies $T_R = 35$ and $T = 70$. The survival probabilities are a decreasing function of age, following the *National Center for Health Statistics life tables* for 1989-1991. We use the age-specific mortality rates for the whole population.

3.2 Preferences

Our utility function has three parameters to be set: ϵ, which captures the elasticity of substitution between goods, the coefficient of relative risk aversion, σ, and the weight of nondurable goods, γ. We set σ to 2, as widely used in the literature. This is the We also fix $\epsilon = 0$, implying a Cobb-Douglas aggregator for durables and nondurables. We do this on empirical grounds, following the evolution of expenditure shares over time in the US. The relative price of durable goods to non-durable goods has fallen steadily at least for the last 40 years, whereas the share of durables to nondurables has remained stable over time, a feature consistent with a unit elasticity of substitution between goods. Consistently with the time series evolution of aggregate data, the estimates of the elasticity of substitution between durables and nondurables from micro data often can't reject the hypothesis of $\epsilon = 0$.[6] The remaining parameter, γ, is calibrated to match the expenditure shares in the data, as detailed below.

3.3 Income process

We calibrate the earnings process to replicate the main features of earnings dispersion in the PSID data.[7] The deterministic component is set to mimic the average age profile of

[5] In principle, education could also be excluded from non-durable consumption, as it can constitue a form of investment rather than consumption. We kept it in our measure of consumption to use a consistent definition across datasets, since we observed it separately in the annual CEX sample, but not in the quarterly panel taken from Heathcote, Perri, and Violante (2010).

[6] See Fernandez-Villaverde and Krueger (2010) for a summary of empirical estimates of this parameter.

[7] We work with the variance of *residual* earnings, which are obtained from a regression of log earnings on a number of controls, including time and age dummies.

after-tax earnings. Following Kaplan and Violante (2010), we choose the variance of the permanent shock to match the increase in earnings dispersion over the life cycle. We have 2 parameters left to match the level of the variance, $\sigma^2_{z_0}$ and σ^2_{ε}. We fixed the proportion between the variance of transitory shocks and permanent shocks to 5, the one implied by the calibration in Kaplan and Violante (2010). The initial variance of the permanent component of income is then set so as to replicate the variance of dispersion of earnings at 25.

3.4 Technology parameters

The return to savings r is set to 3%. The pension benefits are a concave function of average working life earnings, explicitely capturing the progressivity of the U.S. social security system. The specific bendpoints and replacement rates are taken from Storesletten, Telmer, and Yaron (1999).[8]

3.5 Parameters calibrated in equilibrium

We choose γ, β, α, δ, and θ to match a set of cross-sectional statistics measured from US data. Table 1 summarizes the targets and sources used, as well as the calibration outcomes. In particular, we target the average durables to non-durables expenditure ratio in the Consumer Expenditure Survey (CEX), the average wealth to income ratio in the Survey of Consumer Finance (SCF), the fraction of households with negative net worth in the SCF, the ratio between aggregate expenditures in durables and the aggregate stock of consumer durables reported in *Fixed Reproducible Tangible Wealth*, and the average fraction of durable goods purchases which are self-financed in the CEX.

Both in the data and in the model, wealth is defined as total net worth and income as total (labour and financial) income. Since our model is stationary, we net out the growth rate of the stock of durables from the aggregate expenditures series to compute the I/D ratio in the data.[9]

A key parameter in our analysis is θ, which captures the extent to which durable goods can be used as collateral for borrowing. As discussed above, the rebalancing in the consumption basket of constrained households is decreasing in θ, disappearing when $\theta = 1$. We exploit the information available in the CEX on new loans acquired to purchase

[8] Storesletten, Telmer, and Yaron (1999) report the actual figures in 1993 dollars and the relative values with respect to GNP per capita. Our model generates relative values in line with the latter, using total household income to measure GNP, since we don't model production.

[9] Alternatively, we could use the depreciation series estimated by the BEA in the same source to compute a depreciation-to-stock ratio. Both choices deliver very similar numbers.

vehicles. In particular, for each household with positive expenditures in durable goods, we divide the amount borrowed to purchase vehicles by total expenditures in durable goods. We take the average value of this ratio across households as a measure of the extent to which durable goods are self-financed. This calculation delivers a value of 0.10. The value of θ is determined in equilibrium to match this number, since in the model the amount of debt generated by a purchase of durables is $\theta\frac{(1-\delta)}{1+r}I$.

We also calibrate a version of the model without durable goods. In this case, θ, δ, and γ are absent. However, β and α need to be recalibrated in order to keep wealth profiles tied to the data.

Table 1: Calibration targets and results

Parameter	Value	Target	Model	Data
(a) Common				
σ	2			
r	3%			
σ_ε^2	0.05			
σ_η^2	0.01			
$\sigma_{z_0}^2$	0.15			
(b) With durables				
β	0.9518	average W/Y (SCF)	2.5000	2.5000
α	1.0210	fraction with $W<0$ (SCF)	0.1200	0.1200
γ	0.7846	average I/C (CEX)	0.2444	0.2444
θ	0.1252	average $\theta\frac{(1-\delta)}{1+r}$ (CEX)	0.1000	0.1000
δ	0.1775	aggregate I/D (BEA)	0.1891	0.1891
(c) Without durables				
β	0.9584	average W/Y (SCF)	2.5000	2.5000
α	1.0520	fraction with $W<0$ (SCF)	0.1200	0.1200

3.6 Simulated life cycle profiles

Figure 1 shows the average life cycle profile for the main variables in our model, expressed in tens of thousands of dollars. Blue lines depict the profiles emerging from the model without durable goods, while the red lines represent our baseline model with durables. In the top left panel, we can see the characteristic hump shape of consumption expenditures, as documented in Fernandez-Villaverde and Krueger (2006). The top right panel shows a much smaller hump in expenditures in durable goods. It also illustrates the incentive to accumulate durable goods early in life, with high levels of expenditures at early ages. The bottom panels show the evolution of wealth and its composition over the life cycle. Total

Figure 1: Average Life Cycle Profiles

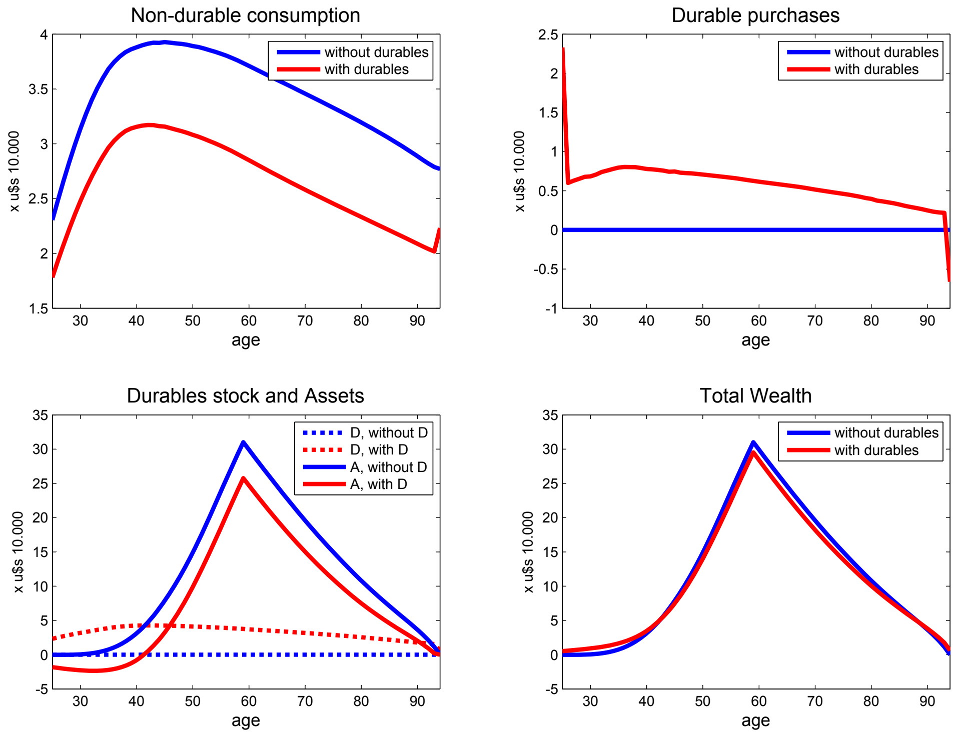

household wealth is shown in the bottom right panel, where it can be seen that households start to accumulate wealth slightly later in life if durables are omitted. However, as the bottom left panel shows, the biggest difference between the two models is in the composition of wealth. Households build up a stock of durables at early ages (dotted red line), long before the accumulation of financial assets begin in the model without durables (solid blue line). In contrast, the average level of financial assets held by households is considerably lower in the model with durables (solid red line), and it is in fact negative during a large portion of the working life, as households get indebted to finance their initial acquisitions of durables.

The life-cycle profiles generated by the model are roughly in line with previous findings in the literature. In particular, nondurable consumption peaks around 45 years of age, as documented by Gourinchas and Parker (2002) and Fernandez-Villaverde and Krueger (2006).[10] Also in line with Gourinchas and Parker (2002), consumption anticipates the

[10]Fernandez-Villaverde and Krueger (2006) find, however, that the timing of the hump in consumption is sensitive to the equivalence scale chosen to compare households of different sizes.

peak in income (not shown), which happens at around 55. However, the size of the hump in non-durable consumption exceeds the empirical counterparts. The increase of nondurable consumption between age 25 and the peak is 70%, compared to the estimated 25% in Fernandez-Villaverde and Krueger (2006) or 50% in Gourinchas and Parker (2002). A similar pattern is observed for expenditures in durable goods, apart from an initial spike in durables expenditures, as our households are born without any durables. This feature is also present in Fernandez-Villaverde and Krueger (2010), who build a very similar model. After the initial period, expenditures in durables are hump-shaped, resembling the empirical findings of Fernandez-Villaverde and Krueger (2006) in terms of size and timing, the peak being approximately 33% higher than the minimum. Total expenditures inherits the excessive hump of nondurables: the increase in total expenditures is 55% in the model, compared to the 30% in estimated by Gourinchas and Parker (2002) and Fernandez-Villaverde and Krueger (2006). Wealth accumulation follows the characteristic pattern in this family of models, with most of the asset accumulation taking place near retirement. In particular, the evolution of wealth in the model without durable goods is comparable to the zero borrowing limit case in Kaplan and Violante (2010)[11], while the wealth composition in the economy with durables shares the main features of Fernandez-Villaverde and Krueger (2010).

Figure 2 illustrates the importance of borrowing constraints over the life cycle. This is important to our analysis, since the effect that we are trying to measure arises when borrowing constraints are binding. Specifically, Figure 2 depicts the cross-sectional average by age of the ex-post discounted ratio of marginal utilities, $\beta\pi_{t,t+1}(1+r_{t+1})\frac{\mu_{t+1}}{\mu_t}$. This is just an ex-post version of the Euler equation, and should be equal to 1 in the absence of binding borrowing constraints. We can see that largest deviations from the optimal intertemporal allocation of consumption are concentrated among young households. By the age of 45, the borrowing constraint is not binding for most households, and, on average, no significant deviation from the desired allocations are observed.

4 The transmission of shocks to consumption and the changes in marginal utility

Let c_{it} be log non durable consumption for household i at age t and d_{it} be log durable consumption for household i at age t. We define the transmission coefficients for the shock

[11] Although we allow for some uncollateralized lending, our specification is overall restrictive in the absence of durables.

Figure 2: Intertemporal allocations

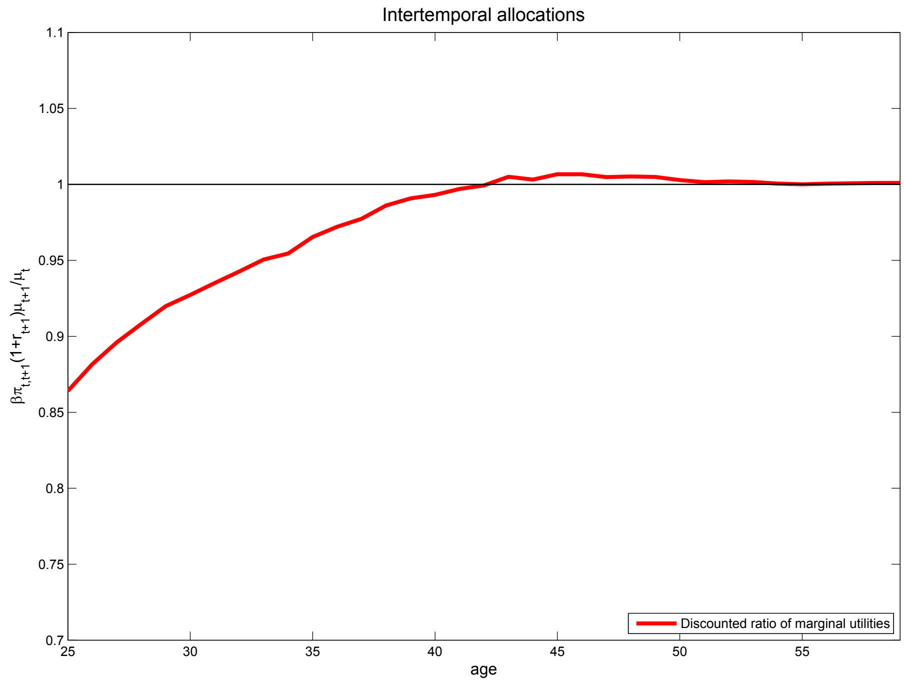

x_{it} as

$$\phi_x^c = \frac{cov\left(\Delta c_{it}, x_{it}\right)}{var\left(x_{it}\right)}$$
$$\phi_x^d = \frac{cov\left(\Delta d_{it}, x_{it}\right)}{var\left(x_{it}\right)}$$

These coefficients measure the proportional change in each consumption good that comes as a response to shocks. ϕ_x^c has been used as a measure of (lack of) insurance because it measures the change in marginal utility of consumption —expressed back in consumption units— in a model without durables. In particular, if the utility function was given by $u\left(C\right) = \frac{C^{1-\sigma}-1}{1-\sigma}$, then

$$\frac{cov\left(\Delta \log\left(\frac{\partial u(C_{it})}{\partial C_{it}}\right)^{-\sigma}, x_{it}\right)}{var\left(x_{it}\right)} = \frac{cov\left(\Delta c_{it}, x_{it}\right)}{var\left(x_{it}\right)} = \phi_x^c$$

To come out with the right measure of insurance in a model with durable goods we also need to quantify the change in the marginal utility of consumption upon arrival of an income shock. In particular we look at the change in the marginal utility expressed in non-durable consumption units of a model without durable goods

$$\phi_x^{uc} = \frac{cov\left(\Delta \log \left(\frac{\partial u(C_{it},D_{it})}{\partial C_{it}}\right)^{-\sigma}, x_{it}\right)}{var\left(x_{it}\right)} = \phi_x^c + \frac{\left(1-\gamma\right)\left(\sigma-1\right)}{\sigma}\left(\phi_x^c - \phi_x^d\right)$$

Notice that the difference of the transmission coefficients $\phi_x^c - \phi_x^d$ gives us the amount of rebalancing between durable and non-durable goods as a response to a shock

$$\phi_x^c - \phi_x^d = \frac{cov\left(\Delta c_{it} - \Delta d_{it}, x_{it}\right)}{var\left(x_{it}\right)} = \frac{cov\Big(\left(c_{it} - d_{it}\right) - \left(c_{it-1} - d_{it-1}\right), x_{it}\Big)}{var\left(x_{it}\right)} \tag{21}$$

Equation (20) shows that in absence of binding borrowing constraints $c_{it} - d_{it}$ is constant over time and independent from shocks. Hence, $\phi_x^c - \phi_x^d = 0$ for both shocks and $\phi_x^{uc} = \phi_x^c$. In that case, the transmission coefficient of non-durable consumption is a correct measure of (lack of) insurance. Instead, when the borrowing constraints are binding, equation (18) shows that $\phi_x^c - \phi_x^d > 0$ if the shock x_{it} makes the borrowing constraints more stringent and $\phi_x^c - \phi_x^d < 0$ if the shock alleviates the borrowing constraints. In this situation, ϕ_x^c gives a biased measure of insurance and the difference $\phi_x^c - \phi_x^{uc}$ tells us how much of the transmission of income shocks into nondurable consumption is due to rebalancing.

Finally, we also define a transmission coefficient for the basket of goods $V = C^\gamma D^{1-\gamma}$,

$$\phi_x^v = \frac{cov\left(\Delta v_{it}, x_{it}\right)}{var\left(x_{it}\right)} = \gamma \phi_x^c + \left(1-\gamma\right)\phi_x^d$$

These transmission coefficients will tell us how much the composite basket of durable and non-durable goods change as a response to shocks.

4.1 Transmission and insurance

Now, we use our calibrated model to compute the transmission coefficient by age with simulated data. In Figure 3 we plot the transmission coefficients for both non-durable goods. The solid line corresponds to the permanent sock and the dashed thick line to the transitory shock. We also plot the transmission coefficients for durable goods (lines with diamonds). The difference between these two lines gives the extent of rebalancing.

Several features arise from the transmission coefficient ϕ_x^c of the non-durable good in Figure 3. In line with the previous literature, permanent shocks have a much larger impact on consumption than transitory shocks. As it is usual in this kind of models, there is a clear life-cycle pattern in the transmission of shocks. For both types of income shocks, the transmission to non-durable expenditures decreases with age. This pattern is qualitatively consistent with the findings of Cerletti (2011), who shows falling transmissions of shocks to consumption for Spanish households. To the best of our knowledge, no similar profile has been documented empirically for the US.[12] The shape of the age profile for transmission is the result of two forces: the age profile of binding borrowing constraints and the closeness to the retirement age. The fraction of households hitting the borrowing limit is higher at young ages, when the accumulated wealth is low. Older households are better self-insured against transitory shocks, explaining the reduction in ϕ^c_ε as age increases. On the other hand, permanent shocks are only permanent in the sense of lasting for the whole working life. Hence, as the retirement age approaches, permanent and transitory shocks are more alike. Therefore, the gap between the transmission coefficients of both types of shocks disappears as households grow old.

Figure 3: Transmission coefficients of income to consumption

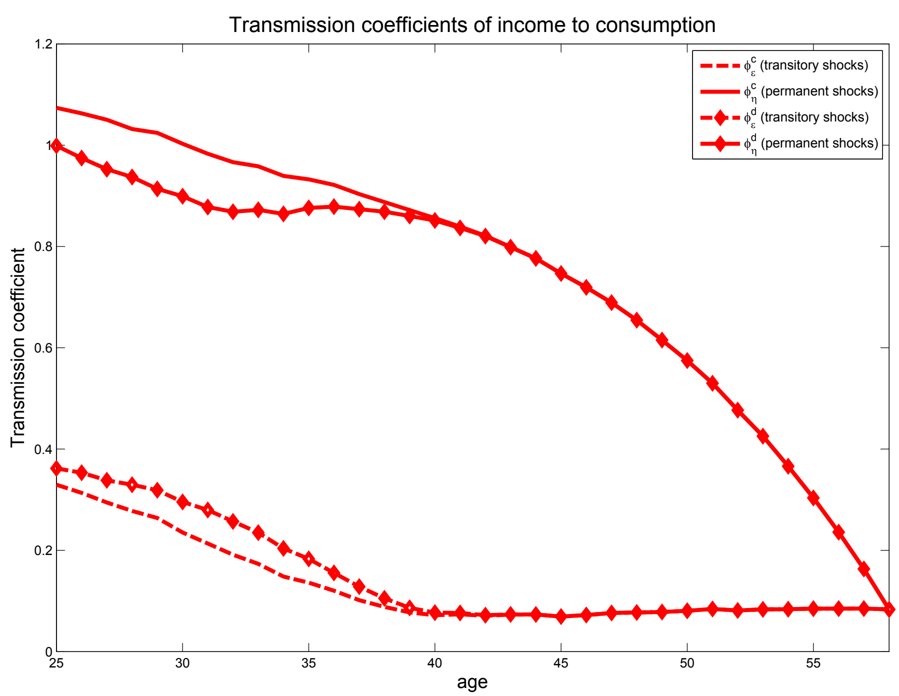

Comparing the transmission coefficients of nondurable consumption with the ones of

[12] Blundell, Pistaferri, and Preston (2008) estimate transmission coefficients for two different cohorts. They obtain mildly higher transmission for the younger cohort, specially with respect to permanent shocks, but the difference across cohorts is not statistically significant.

the durable good, we find important differences for young (constrained) households. In particular, we can see that the fraction of a transitory shock that passes on to non-durable consumption is lower than its equivalent for durables. On the other hand, a permanent shock has a higher impact over non-durable consumption than it has over durable goods. These differences reflect the consumption rebalancing described in section 2.9.1, and they tell us that the transmission of income shocks to nondurable consumption cannot be interpreted as a measure of insurance in young households. In particular, the amount of insurance against transitory shocks is lower than what the transmission coefficients suggest, while the amount of insurance against permanent shocks is higher than what the transmission coefficients suggest.

Another way of seeing this same results is in Figure 4, where we plot the amount of rebalancing computed as in equation (21). We can see that consumption rebalancing is important for young households, disappearing after the age of 45, when liquidity constraints cease to bind, as seen in Figure 2. A transitory shock to a constrained household induces a rebalancing towards durable goods ($\phi^c_\varepsilon < \phi^d_\varepsilon$), hence the response of non-durables is lower than that of the consumption basket. A permanent shock to a constrained household has the opposite rebalancing effect, rebalancing towards nondurable goods ($\phi^c_\eta > \phi^d_\eta$), and hence transmitting more to non-durable goods than to the composite basket. In terms of equation (20), this implies that λ_t/μ_t, our measure of the tightness of the borrowing constraint, comoves positively with the permanent shock. In other words, positive permanent shocks aggravate the severity of borrowing constraints, while negative permanent shocks ease it. These differences disappear over the life cycle, as borrowing constraints become less binding on average, and the responses to shocks of both goods converge.

Finally, we can look at the average of the transmission coefficients. Table 2 summarizes the results on transmission of income shocks into the different consumption goods. Panel (a) reports the coefficients for all households, panel (b) shows the same information for the youngest households only, and panel (c) for the oldest (working-age) households. Again, we can see that consumption is much more sensitive to income shocks for young households, which are more likely to be constrained, than for old ones. In fact, non-durable consumption of the young reacts practically one to one with permanent shocks, while only 58% of permanent shocks are transmitted into non-durable consumption for the old. Overall, the average transmission of permanent shocks for all households is 74%. For transitory shocks, we find that 13% are transmitted into non-durable consumption for the young and slightly less than 8% for the old, with an average transmission over the whole population of 13%. In comparison to the empirical findings by Blundell, Pistaferri, and Preston (2008), our model generates larger transmission of shocks to nondurable

Figure 4: Difference in transmission coefficients

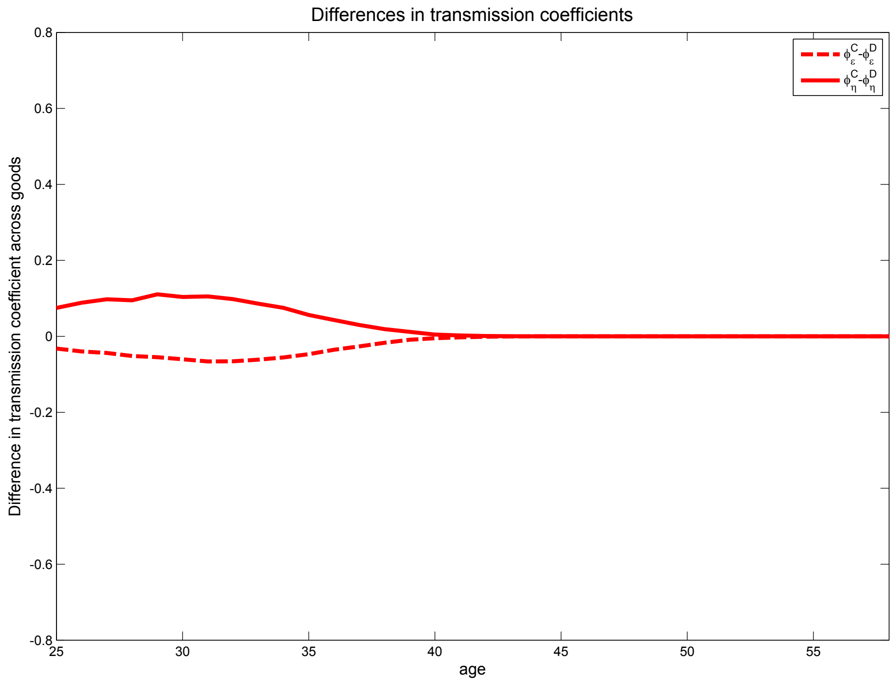

consumption, roughly 10 percentage points higher than their baseline estimates for both types of shock. This could imply that our model economy lacks some smoothing mechanism available to the households in their sample, a point the literature has made in the context of single-good models.[13] Perhaps remarkably, our transmission coefficients are in line with those obtained by Blundell, Pistaferri, and Preston (2008) once they augment their baseline sample with a subsample of low income households.[14] If these additional households face tighter borrowing constraints than their representative sample, this would point to our specification of the borrowing limit being excessively tight at least for some households.

Regarding the transmission of income shocks into durable goods, they are similar to the ones for non-durable goods, but we find some differences among young households. In particular, for young households, the transmission on durable goods is 5 percentage

[13]See Kaplan and Violante (2010) for a discussion on the limitations of self-insurance as the only mechanism to smooth consumption, and Kaplan (2009) for a discussions of potential insurance channels for the young.

[14]They exclude the SEO sub-sample of PSID for most of their analysis, which comprises households who had low-income in 1968 and their subsequent split-offs. When included, they obtain a transmission coefficient of 12.1% for transitory shocks and 76.5% for permanent shocks.

Table 2: Transmission coefficients of income to consumption

	Nondurables (C) ϕ_x^c	Durables (D) ϕ_x^c	Bias $\phi_x^c - \phi_x^{uc}$
(a) All households			
Transitory shocks	13.1	15.1	-0.2
Permanent shock	74.3	71.0	0.4
(b) Young households (below 40)			
Transitory shocks	20.6	25.3	-0.5
Permanent shock	98.0	90.3	0.8
(c) Old households (over 40)			
Transitory shocks	7.8	7.9	0.0
Permanent shock	57.8	57.7	0.0

Note: transmission coefficients are expressed in percentage points.

points higher for the transitory shocks and 8 percentage points lower for the permanent shock. The third column in Table 2 reports the bias due to rebalancing. Instead, the differences of transmission among the old are almost null.

These differences in transmission imply a very small bias of the insurance measure based on non-durable consumption. As the third column in Table 2 shows, the bias of the transmission of permanent shocks for the young is 0.8 percentage points. That is to say, the transmission of permanent shocks into non-durable consumption is 98.0 percent, and only 0.8 percentage points reflect rebalancing whereas the remaining 97.2 reflect lack of insurance. For transitory shocks, the transmission is 20.6 percentage points and without rebalancing this would only be 21.1, with the bias being -0.5 percentage points. Note that the bias, which we define as the difference between ϕ_x^c and ϕ_x^{uc}, is given by

$$\frac{(1-\gamma)(\sigma-1)}{\sigma}\left(\phi_x^d - \phi_x^c\right)$$

so, with the calibrated values of $\gamma = 0.8$ and $\sigma = 2$, the bias is 0.1 times the difference between the transmission coefficients.

4.2 Differences between models

A final exercise is to compare the transmission coefficients predicted by our model to those of the model without durable goods. The reason for this is that the two models do not differ only in the substitution between goods, but also in some other important aspects. First, the model with durables has a higher ability to borrow (as long as $\theta \neq 0$).

Second, the two economies are calibrated to the same total wealth and the same fraction of constrained households. But as seen in Figure 1, the timing of the wealth accumulation is different: in the model with durables households accumulate more wealth in the first part of the life-cycle, whereas in the model without durables households have more wealth in the second part of the life cycle. Moreover, the composition of wealth is also different: in the model with durable goods, most of the assets held before the age of 50 are durable goods, and financial assets are in fact negative for the first 15 years. As reported in table 1, households are somewhat more impatient in the model with durables.

In Figure 5 we plot the transmission coefficients for nondurable consumption for both models. Red lines depict the transmission coefficients in the model with durable goods, while blue lines indicate the transmission coefficients in the model without durables. The solid lines correspond to the transmission of permanent shocks, and the dashed lines to the transitory shock. We see that the main differences are in permanent shocks and in younger households. In particular, the transmission of permanent shocks to nondurable consumption is lower in the model with durables up to age 40, being slightly higher thereafter. In contrast, the transmission of transitory shocks is slightly higher in the model with durables at young ages, with all the differences disappearing in the second half of the working life. The differences of transmissions by age groups below 40 go in the opposite direction than rebalacning alone would suggest. Since households are more patient in the model without durables, they keep higher liquid assets to insure themselves against transitory income fluctuations. On the other hand, the transmission of permanent shocks reflect the trade-off that constrained households face: at the constraint, there is a unique mapping from the level of non-durable consumption (C_t) to the composition of consumption (C_t/D_t). Hence, they optimally reduce the overall response of non-durable consumption to avoid departing too far from the unconstrained consumption bundle.

5 Extensions and robustness checks

Although our model assumptions and calibration targets are empirically motivated, as described above, some of them are by no means incontrovertible. In this section, we analyze the sensitivity of our findings to alternative modeling choices and calibration strategies. Table 3 summarizes the calibration results of the different exercises we conducted. We explain each of them in detail below.

Figure 5: Transmission coefficients (non-durables)

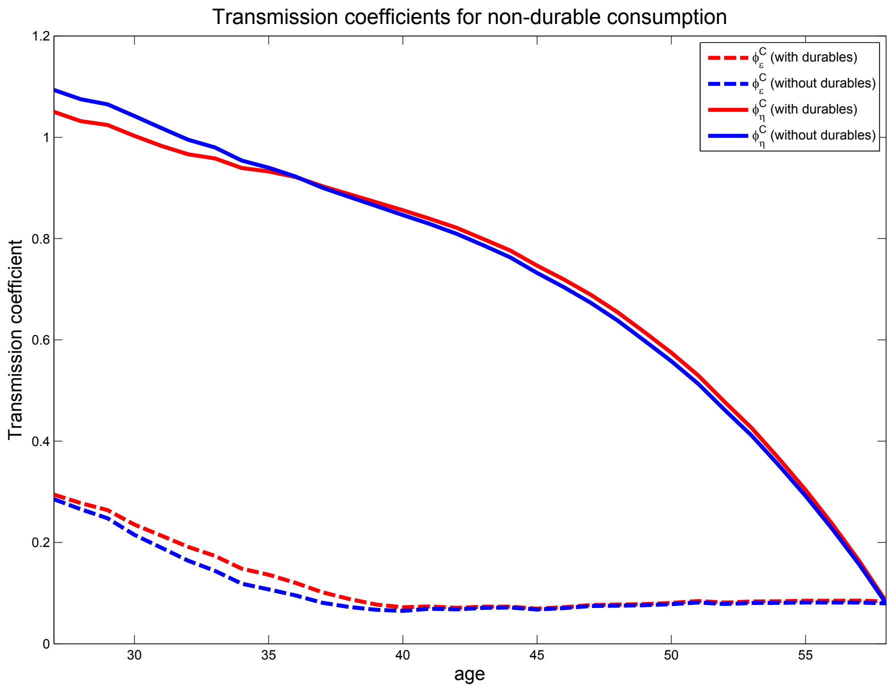

5.1 Durability of goods

Equation (18) states that a high durability (a low depreciation rate) of durable goods increases the extent of rebalancing. In our baseline calibration, we obtain a value for δ of 17.86%, which is somewhat higher than the values used in previous studies. For instance, within a very similar framework Aaronson, Agarwal, and French (2008) use a quarterly depreciation rate of 3.4%, which amounts to an annual rate of 12.92%. This is the value reported by Campbell and Hercowitz (2003), who use our same data source to compute separate depreciation rates for housing and other consumer durables.[15]

As a first robustness exercise, we recalibrat our model to match the same targets (except for the ratio of expenditures to stock for durables) while fixing δ to the Campbell and Hercowitz (2003) value. The results of this calibration are summarized in the second column of Table 3. Fixing δ induces little change in the other parameters. However, the model misses the I/D ratio, given the almost direct link between this target and the depreciation rate. Figure 6 shows no sign of significant changes in the timing of wealth

[15] The depreciation rate of housing is much lower than that of consumer durables. Hence, our δ is not directly comparable to the numbers used in the housing literature.

Table 3: Alternative parameterizations: calibration results.

Specification	Baseline	Lower depreciation	No unsecured debt	Illiquid wealth
Parameters				
β	0.9518	0.9501	0.9324	0.9073
α	1.0210	1.0390	0.0000	0.0000
γ	0.7846	0.7755	0.7853	0.7899
θ	0.1252	0.1183	0.1253	0.1253
δ	0.1775	0.1292	0.1790	0.1783
τ	-	-	-	0.0804
Statistics				
average W/Y	2.5000	2.5000	2.5000	0.9073
fraction with $W < 0$	0.1200	0.1200	0.0000	0.0000
average I/C	0.2444	0.2444	0.2444	0.2444
aggregate I/D	0.1891	0.1410	0.1891	0.1891
average W^{liq}/W^{noliq}	1.514	0.8726	1.8560	0.2279

accumulation, which is reassuring in the sense of δ not playing an important role in the age composition of constrained households.

Figure 7 illustrates the effect on rebalancing of decreasing δ. There is a sizable increase in the extent of the rebalancing, more so for permanent shocks. These results are straightforward given the amplifying role of $(1 - \delta)$ in equation (18).

We find that the overall transmission of shocks to nondurable consumption is not affected very much, but its interpretation is. In the second column of Table 4 we report the average transmission coefficients for this economy, which can be easily compared to the ones of the baseline calibration reproduced for convenience in column 1. We can see that the size of the bias doubles for the young (and for the whole population) for both shocks. Hence, under this calibration, the transmission of permanent shocks into non-durable consumption for the young is 96.2% but the transmission into marginal utility would be of only 94.3%, with the difference of 1.9% being accounted by the rebalancing between goods. For the transitory shocks, the transmission into non-durable consumption is of 20.2%, but 1% is due to rebalancing, so the transmission into marginal utility would be 21.2%.

5.2 Unsecured borrowing

In our baseline model, we allow for both collateralized and uncollateralized debt, with θ and α, respectively, measuring the tightness of the borrowing limit with respect to each type of debt. However, most of the literature on durable goods preclude unsecured

Table 4: Alternative parameterizations: transmission of income shocks.

Specification	Baseline	Lower depreciation	No unsecured debt	Illiquid wealth
Transmission (all)				
ϕ_ε^c	13.1	13.0	19.8	26.7
ϕ_η^c	74.3	73.6	79.0	86.8
ϕ_ε^d	15.1	16.6	23.9	34.0
ϕ_η^d	71.1	66.5	72.9	76.8
$\phi_\varepsilon^c - \phi_\varepsilon^{uc}$	-0.2	-0.4	-0.4	-0.7
$\phi_\eta^c - \phi_\eta^{uc}$	0.4	0.8	0.6	1.1
Transmission (young)				
ϕ_ε^c	20.6	20.2	32.0	35.9
ϕ_η^c	98.1	96.2	103.6	104.7
ϕ_ε^d	25.3	28.7	39.8	44.1
ϕ_η^d	90.3	79.5	91.1	90.5
$\phi_\varepsilon^c - \phi_\varepsilon^{uc}$	-0.5	-1.0	-0.8	-0.9
$\phi_\eta^c - \phi_\eta^{uc}$	0.8	1.9	1.3	1.5
Transmission (old)				
ϕ_ε^c	7.8	7.9	11.2	20.3
ϕ_η^c	57.8	57.9	62.1	74.5
ϕ_ε^d	7.9	8.2	12.8	26.9
ϕ_η^d	57.7	57.6	60.4	67.5
$\phi_\varepsilon^c - \phi_\varepsilon^{uc}$	0.0	0.0	-0.2	-0.7
$\phi_\eta^c - \phi_\eta^{uc}$	0.0	0.0	0.2	0.7

Note: transmission coefficients are expressed in percentage points.

Figure 6: Average Life Cycle Profiles

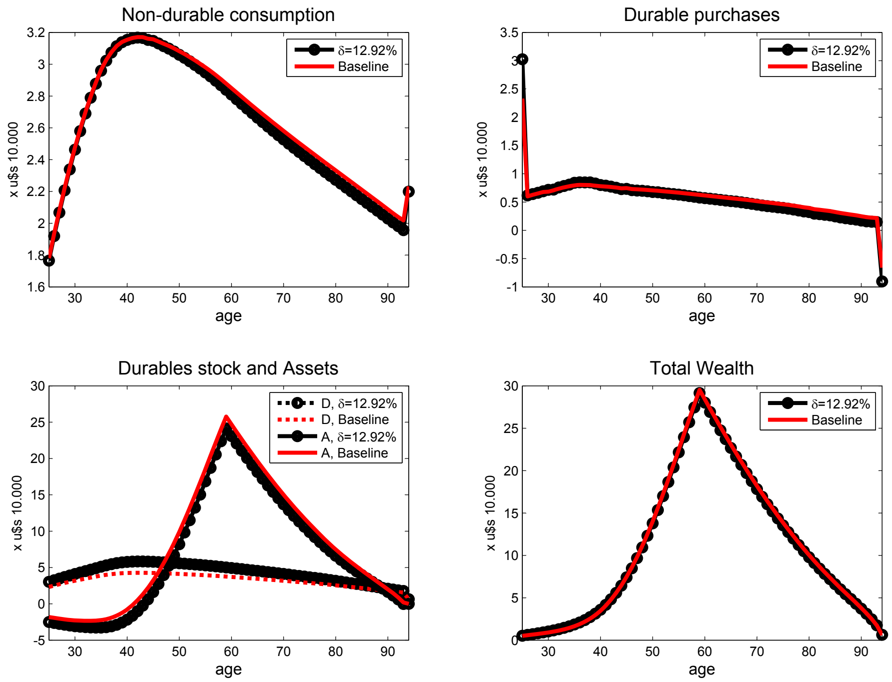

borrowing altogether. This would be the natural outcome from a limited enforceability problem where the seizure of tangible assets is the only punishment for defaulting households.[16] Our definition of the borrowing limit $\underline{A}$ is very restrictive as a consequence of the assumed earnings process, which allows for realization arbitrarily close to zero. However, since our model economy features no risk after retirement, the borrowing limit increases as retirement approaches.

In order to check the importance of unsecured borrowing for the life-cycle profile of binding constraints and rebalancing effects, we compute a version of the model with $\alpha = 0$. In this case, we no longer can generate households with negative net worth, but we recalibrate the model to match the wealth to income ratio, the share of durable goods in expenditures, and the ratio of expenditures to stock of durable goods. The third column of table 3 shows the results of this calibration. The taste for nondurable goods and the depreciation rate are unaffected by the change in α. However, the tighter borrowing limit induces additional precautionary savings with respect to our benchmark. Hence, a lower discount factor is needed to match the observed level of wealth. Figure 8 shows the role of unsecured debt over the life cycle. Nondurable consumption grows faster and

[16]See section 2.4 for a discussion on the parametric restrictions implied by such an assumption.

Figure 7: Difference in transmission coefficients

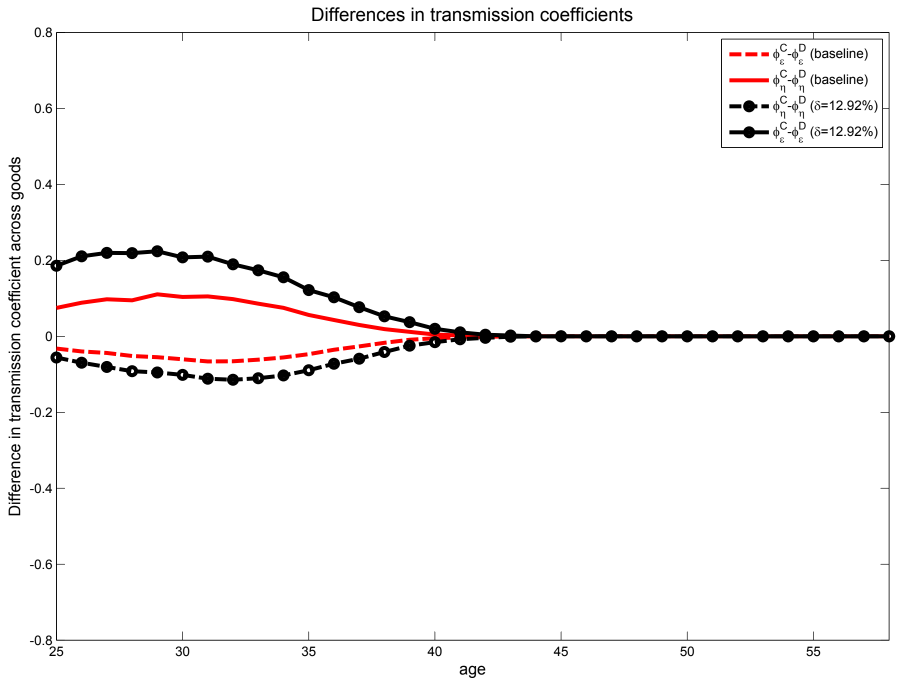

exhibits a larger hump. As a consequence of the lower discount factor, it also decreases more rapidly after its peak. At the same time, driving α to zero removes most of the initial spike in expenditures in durables, spreading the creation of a stock of durables over several periods. Overall, total expenditures are lower early in life, since they are harder to finance. Restricting debt also changes the composition of wealth over the life cycle: while the average stock of durables is essentially the same, the average holdings of financial assets shifts towards younger households. This shift is a combination of three forces: first, precluding unsecured borrowing mechanically increase the net worth of constrained households; second, for given preferences and income risk, a tighter borrowing limit induces higher precautionary savings; and third, the discount factor required to match a given wealth to income target is lower when $\alpha = 0$, effectively compressing the wealth distribution.

We find that when unsecured debt is unavailable, the rebalancing effects are larger and are present over a longer period of life. Figure (9) shows that, for both types of shocks, rebalancing is larger for young people and it remains different from zero 5 years longer than in the baseline calibration. As discussed above, the availability of unsecured debt increases with age until retirement, contributing to the marked age profile in transmission

Figure 8: Average Life Cycle Profiles

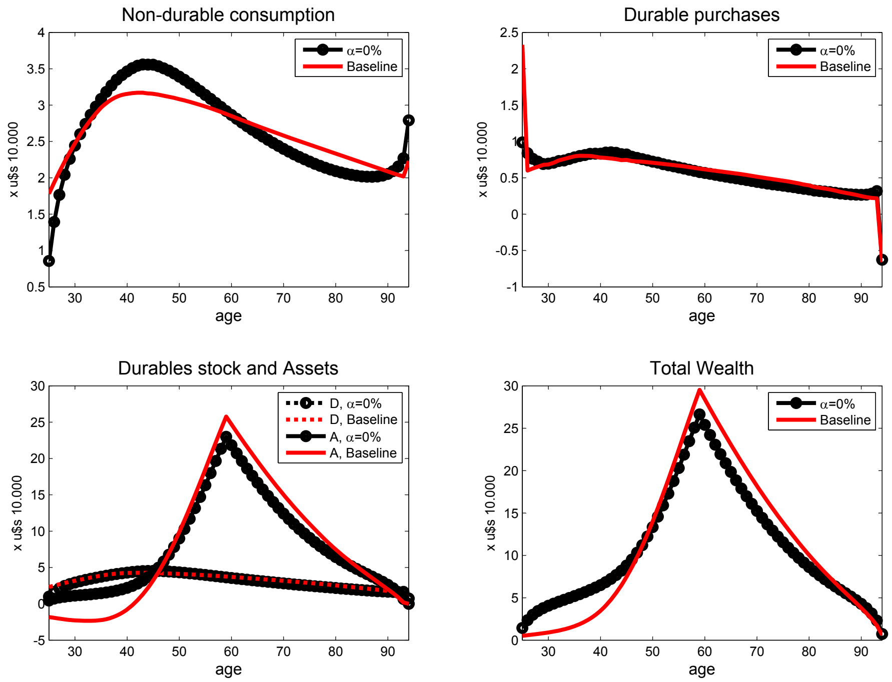

of income shocks to consumption. Once unsecured debt is removed, the overall ability to smooth shocks is lower, but more so for older households. The third column in Table 4 shows how rebalancing translates into the difference between transmission to nondurable consumption and insurance. The overall bias increases from -0.2 to -0.4 in the case of transitory shocks, and from 0.4 to 0.6 in the case of permanent shocks. A large part of the increase in the bias come form young households. However, in the absence of unsecured debt, the bias also becomes non-zero for households over 40 years old. In particular, it is as big as the overall bias in the benchmark case for transitory shocks, and half the overall benchmark bias for permanent shock.

5.3 Liquidity of assets

So far, we have considered a single financial asset, A_t, to reflect the total net worth of households, once consumer durables are excluded. This definition hides the heterogeneous nature of the different components of households' balance sheets, specially in terms of liquidity. Moreover, a fraction as large as 80% of average household wealth is held in illiquid assets. Therefore, the self-insurance role of wealth may be overstated in our

Figure 9: Difference in transmission coefficients

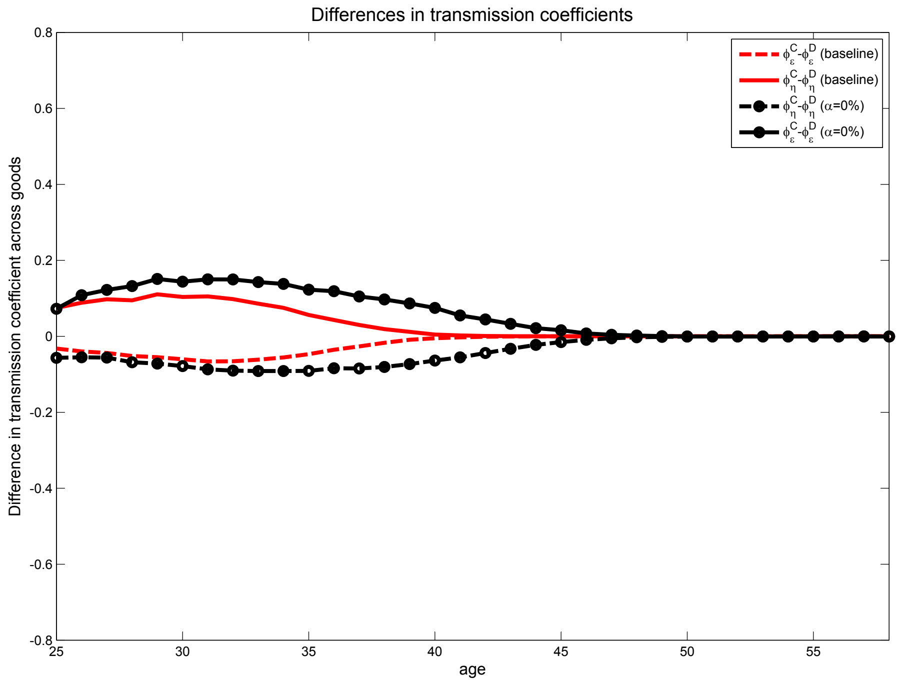

baseline exercise. Modelling portfolio decisions in the presence of illiquid assets is beyond the scope of this paper. However, we acknowledge that our characterization of A_t can be interpreted as an extreme assumption about portfolio composition, where all wealth is held in the liquid asset.

Hence, for completeness, we run a simple alternative extreme case, in which some fraction of total wealth is fully illiquid and can not be accessed before retirement. The interpretation is that households save in illiquid assets for life-cycle considerations, while they keep liquid assets for precautionary motives. Specifically, we maintain liquid savings as an endogenous variable, but we restrict savings in illiquid assets, which we label "retirement accounts", to be a constant fraction τ of income. Upon retirement, households receive the (capitalized) value of retirement accounts as a lump-sum transfer.[17] Therefore, the baseline model assumes that the retirement accounts can be withdrawn in full at any time and no cost, while the alternative forbids any anticipated withdrawal. To be consistent with the definition used in the data, we compute net liquid (illiquid) wealth

[17]Since life is deterministic after retirement, recovering illiquid assets as wither a one-time payment or an annuity does not matter. By abstracting from any potential misalignment between the timing of illiquid assets payments and the desired consumption profile, the one-time assumption ensures that after-retirement life is the same in both the baseline and the alternative model.

Figure 10: Average Life Cycle Profiles

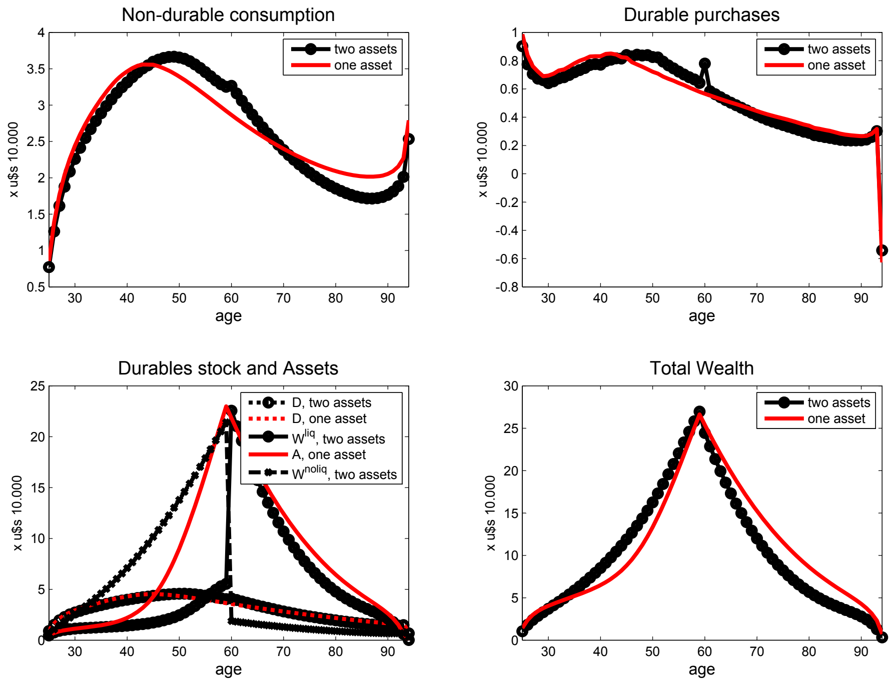

as total liquid (illiquid) assets minus total liabilities associated to the purchase of liquid (illiquid) assets. In terms of our model, we can define liquid wealth W_t^{liq}, and illiquid wealth W_t^{noliq}, as:

$$
\begin{aligned}
W_t^{liq} &= A_t + \theta \frac{1-\delta}{1+r_{t+1}} D_t \\
W_t^{noliq} &= (1 - \theta \frac{1-\delta}{1+r_{t+1}}) D_t + \text{retirement accounts}
\end{aligned}
$$

In table 3 we apply the same definition to single-asset economies, where retirement accounts are not present. Hence, in those cases the ratio W^{liq}/W^{noliq} is just a measure of the contribution of the next stock of consumer durables to total net worth.

Introducing illiquid wealth, even in the simple way outlined above, requires a modification of the calibration strategy. We can still use total wealth to compute the wealth to income ratio as before, but now we also need to pin down the composition of wealth, which is governed by τ. We calibrate β, γ, δ, and τ jointly to match the same statistics as in the previous section plus the liquid to illiquid wealth ratio for working age households. Since

W_t^{noliq} captures the net value of illiquid assets (with the exception of consumer durables), it is unclear how the borrowing limit should be specified in the two-asset case. We choose to maintain $\alpha = 0$, so the results are directly comparable to the ones obtained in Section 5.2. The calibration of this two-asset economy is presented in the last column of table 3. Two things are worth noticing in Table 3. First, the value of τ obtained resembles the fraction of lifetime earnings held as wealth at retirement. Using the PSID, Hendricks (2007) estimates the average lifetime earnings of a household by computing the capitalized sum of earnings over the life cycle, and finds that average wealth at retirement amounts to 8% of his measure.[18] Second, the discount factor is lower than in the single asset case (third column), as a consequence of the timing assumed for illiquid wealth accumulation. Figure 10 shows that adding the illiquid assets leads to an anticipation of wealth accumulation. This is a result of savings in retirement accounts being proportional to income throughout the working life, whereas in the single-asset economy, savings for retirement are concentrated towards the end of working life.[19] However, liquid wealth held by workers is lower in the two-asset economy, which translates into less self-insurance. Remarkably, the main differences in liquid wealth accumulation between the two economies arise after the age of 40, when life-cycle motives for saving become more important.

Figure 11 shows that, early in life, rebalancing effects are equally important in the two economies, but they disappear much later in the two-asset case, leading to a higher overall incidence of rebalancing and a flatter age profile for transmission coefficients. The coefficients generated by the two-asset model are reported in the fourth column of table 4. Consumption responses are in general larger in this version of the model, since self-insurance is restricted to only a fraction of total wealth. Compared to the model with no unsecured debt, the transmission of transitory shocks into non-durable consumption increases from 19.8% to 26.7%, whereas the transmission of permanent shocks increases from 79.0% to 86.8%. A closer examination of panels (b) and (c) of table 4 reveals that most of the differences come from older households: while the differences in transmission coefficients are smaller than 4 percentage points for households below 40 years old, they can be as large as 14 percentage points for older households. In particular, the transmission of transitory shocks is twice as large for this age group. The same decomposition by age applies to the bias of transmissions due to rebalancing. The differences are small for young households, while households above 40 drive most of the increase in the overall

[18]In our model, average total wealth at retirement, which includes liquid assets and durable goods, is 10% of lifetime earnings, computed in the same way as Hendricks (2007).

[19]It is not obvious how the timing of illiquid wealth accumulation may differ from that of liquid wealth when both assets are endogenously chosen. We refer to Kaplan and Violante (2011) for a recent analysis of wealth composition when all assets are chosen endogenously.

Figure 11: Difference in transmission coefficients

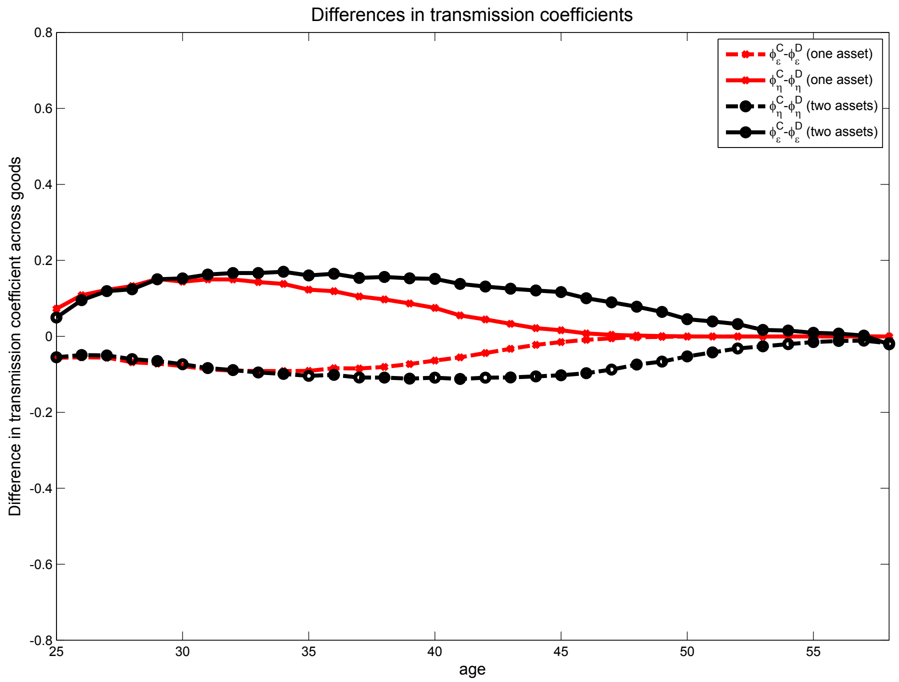

bias.

The conclusions of this exercise are important. If we think that not all household wealth can be used cheaply to accommodate unexpected income changes, a standard life-cycle model of consumption predicts much less insurance than measured in the data. Hence, the excess smoothness puzzle could be severe.

6 Conclusions

In this paper, we have analyzed the responses to income shocks of households that care for both durable and non-durable goods and face borrowing constraints. The main purpose of the analysis was twofold. First, we wanted to characterize the specific responses of the consumption of each type of good. Second, we wanted to asses the impact of neglecting durables in measuring consumption insurance. To this end, we have constructed a life-cycle, incomplete markets model with two goods of different durability. We used the model to characterize the consumption responses to income shocks as a function of liquidity restrictions, the persistence of the shocks, and the durability of the goods. Then, we calibrated the model to replicate the US economy in order to measure the quantitative importance of durable goods for measuring the extent of insurance.

Our main qualitative findings can be summarized as follows. First, we have shown that, in the absence of binding borrowing constraints, the consumption of both durable and non-durable goods responds equally to income shocks. This implies that both goods are consumed in the same proportion regardless of the shock. However, when borrowing constraints bind, there is a rebalancing effect that shifts consumption towards one of the goods depending on the persistence of the shock. When the shock is permanent, non-durable consumption reacts more than durable consumption, whereas the opposite is true when the shock is transitory.

Second, we have shown that insurance, defined as the ability to smooth the marginal utility of consumption across states, is a function of the transmission of income shocks to non-durable consumption and the extent of rebalancing. Therefore, the response to shocks of non-durable consumption alone, even if correctly measured, is not an exact measure of insurance for constrained households.

The quantitative results of the calibrated model are the following. First, we found rebalancing effects to be moderate and concentrated at young ages. The latter result is a consequence of liquidity constraints being more important for younger households. Second, the impact of rebalancing on our measure of insurance is small, specially for transitory shocks. In our baseline calibration, the difference between the transmission of shocks to non-durable consumption and insurance was less than half a percentage point. This difference was bigger for permanent shocks and for young households, but always below one percentage point.

We conducted a series of robustness checks that confirmed the limited role of rebalancing in consumption insurance. These exercises delivered some additional results. In particular, the size of the bias caused by measuring insurance as the transmission to non-durable consumption alone depends on the durability of the other good, although the transmission itself does not. We also found that the availability of uncollateralized loans matters for the age distribution of constrained households, and hence for the incidence of rebalancing over the life cycle. Finally, we found that savings' liquidity can potentially play a role both in the level and the age profile of consumption rebalancing as a response to income shocks. A more careful study of the use of illiquid assets and its links with precautionary and life-cycle motives for wealth accumulation would be needed to draw further conclusions on this issue.

References

Aaronson, D., S. Agarwal, and E. French (2008): "The Spending and Debt Response to Minimum Wage Hikes," Federal Reserve Bank of Chicago, WP 2007-23.

Bils, M., and P. Klenow (1998): "Using Consumer Theory to Test Competing Business Cycles Models," *Journal of Political Economy*, 106(2), 233–261.

Blundell, R., L. Pistaferri, and I. Preston (2008): "Consumption Inequality and Partial Insurance," *American Economic Review*, 98(5), 1887–1921.

Browning, M., and T. Crossley (2009): "Shocks, Stocks, And Socks: Smoothing Consumption Over A Temporary Income Loss," *Journal of the European Economic Association*, 7, 1169–1192.

Campbell, J., and A. Deaton (1989): "Why is Consumption so Smooth?," *Review of Economic Studies*, 56(3), 357–374.

Campbell, J., and Z. Hercowitz (2003): "The Dynamics of Work and Debt," Chicago Fed Working Paper 2004-05.

Cerletti, E. A. (2011): "Life-cycle patterns in consumption insurance: evidence from the Spanish Household Budget Survey," Mimeo, CEMFI.

Chah, E. Y., V. A. Ramey, and R. M. Starr (1995): "Liquidity Constaraints and Intertemporal Consumer Optimization: Theory and Evidence from Durable Goods," *Journal of Money Credit and Banking*, 27(1), 272–287.

Fernandez-Villaverde, J., and D. Krueger (2006): "Consumption over the Life Cycle: Facts from the Consumer Expenditure Survey," *Review of Economics and Statistics*, 3(89), 552–565.

——— (2010): "Life-Cycle Consumption, Debt Constraints and Durable Goods," Forthcoming in *Macroeconomic Dynamics*.

Gourinchas, P. O., and J. A. Parker (2002): "Consumption Over the Life-Cycle," *Econometrica*, 70, 47–89.

Hall, R., and F. Mishkin (1982): "The Sensitivity Of Consumption To Transitory Income: Estimates From Panel Data On Households," *Econometrica*, 50, 461–481.

Harris, E., and J. Sabelhaus (2000): "Consumer Expenditure Survey. Family-level Extracts, 1980:1 - 1998:2," Mimeo, Congressional Budget Office.

Heathcote, J., F. Perri, and G. Violante (2010): "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967- 2006," *Review of Economic Dynamics*, 1(13), 15–51.

HENDRICKS, L. (2007): "Retirement Wealth and Lifetime Earnigns," *International Economic Review*, 48(2), 421–456.

JOHNSON, D., R. MCCLELLAND, AND J. PARKER (2010): "Consumer Spending and the Economic Stimulus Payments of 2008," Mimeo.

KAPLAN, G. (2009): "Moving Back Home: Insurance Against Labor Market Risk," Mimeo, NYU.

KAPLAN, G., AND G. VIOLANTE (2010): "How Much Consumption Insurance Beyond Self-Insurance?," 2(4), 53–87.

KAPLAN, G., AND G. L. VIOLANTE (2011): "A Model of the Consumpiton Response to Fiscal Stimulus Payments," manuscript.

LUENGO-PRADO, M. J. (2006): "Durables, Nondurables, Down Payments and Consumption Excesses," *Journal of Monetary Economics*, 53(7), 1509–1539.

STORESLETTEN, K., C. TELMER, AND A. YARON (1999): "The Risk Sharing Implications of Alternative Social Security Arrangements," *Carnegie-Rochester Series on Public Policy*, 50, 213–259.