Durable Goods, Borrowing Constraints, and Consumption Insurance*

Enzo A. Cerletti
Goethe University Frankfurt

Josep Pijoan-Mas
CEMFI and CEPR

May 15, 2014

Abstract

We analyze the transmission of income shocks into durable consumption goods. We show that binding borrowing constraints lead to a substitution between goods of different durability upon arrival of an unexpected income change. The sign of this substitution depends on the persistence of the shock, whereas its size depends on the durability of goods and on their role as collateral for borrowing. An important consequence is that the response of nondurable consumption to income shocks may be an imperfect measure of household insurance against labor market risk. We use a calibrated two-good life-cycle model with labor market uncertainty and incomplete markets to quantify the actual amount of insurance implied by the observed transmission of income shocks to nondurable consumption. We find that young households have substantially less insurance against transitory shocks and more insurance against permanent shocks than commonly thought.

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

*This paper has benefited from comments by assistants to seminars held at Banque de France, CEMFI, Goethe University, Macro Club at UPenn and National Bank of Serbia. It has also benefited from discussions during the presentations at the Annual Congress of the European Economic Association (Oslo), the Annual Meetings of the American Economic Association (Chicago), and the Simposio de la Asociación Española de Economía (Madrid). Enzo Cerletti gratefully acknowledges financial support from the Spanish Ministry of the Economy and Competitiveness (grant ref. BES-2009-019394). Josep Pijoan-Mas gratefully acknowledges financial support from the Spanish Ministry of the Economy and Competitiveness (grant ref. ECO2010-16726). Postal address: CEMFI, Casado del Alisal 5, 28014 Madrid, Spain. E-mail: cerletti@econ.uni-frankfurt.de, pijoan@cemfi.es

1 Introduction

The standard life-cycle model of consumption under complete markets predicts that households smooth out income fluctuations altogether. As a result, at least since Hall and Mishkin (1982), the response of expenditure of nondurable consumption goods to unexpected income changes has been used to measure the amount of insurance available to private households. But the use of nondurable consumption responses to income shocks as a measure of insurance ignores substitution with durable goods. Yet, there is growing evidence that expenditure on durable goods reacts much more to unexpected income changes than expenditure in nondurable goods.[1] Therefore, in order to use consumption data to learn about household insurance, we need first to understand the substitution between durable and nondurable goods upon arrival of unexpected income changes.

In this paper we make a first step in this direction and study the response to income shocks of goods of different durability. In particular, we assume homothetic preferences, we abstract from adjustment costs, and we focus on the interplay of borrowing constraints and the persistence of shocks. Given these assumptions, we think of durable goods as cars, furniture, home appliances, and the like, but we exclude housing from our main exercise.[2] As it is well known, under these assumptions and no binding borrowing constraints, a standard consumption model predicts that the optimal composition of the basket between durable and nondurable goods is given by the user cost of durables, which does not change with income. Hence, the response to income shocks is identical across all goods, and it is determined by the amount of insurance available. However, in the presence of binding liquidity constraints, households shift their consumption basket away from more durable goods because their future services are valued less when households would like to bring consumption to the present and cannot do so.[3] In this context, unexpected income changes affect the severity of the borrowing constraints and the consumption responses to shocks become a mix of the lack of insurance with the substitution between goods of different durability.

Our first contribution is to characterize the different transmission of unexpected income changes into consumption goods of different durability. For clarity, we consider only one durable and one nondurable consumption good. We show that the persistence

[1] See for instance Browning and Crossley (2009), Johnson, McClelland, and Parker (2013) or Aaronson, Agarwal, and French (2012).

[2] Excluding housing from the definition of durable goods is relatively common when studying consumption responses to labor income shocks, see for instance references in footnote 1. Houses are lumpy goods whose value is several times larger than the income shocks we are looking at, so it is an unlikely margin of adjustment.

[3] See Chah, Ramey, and Starr (1995).

of shocks is critical because it affects the marginal propensity to consume out of the income change, and hence the severity of the borrowing constraints. Households that face a transitory positive income shock want to spread the income innovation over their lifetime, which leads them to save most of the income increase. For constrained households, this alleviates the borrowing constraint, which spurs a substitution from nondurable goods towards durable goods, as well as the standard effect of increasing expenditure in both types of goods. As a a consequence, the response of durable goods to the income shock is larger than the response of nondurable goods. This difference depends positively on the durability of the goods and negatively on the fraction of durable good expenditure that can be collateralized. [4] When income shocks are more persistent, households desire to save a lower fraction of the income increase because the shock also contains information about future higher income. For constrained households, this means that more persistent shocks alleviate the borrowing constraints to a lesser extent, and hence imply a smaller substitution between goods and a smaller difference in the responses to shocks of durable and nondurable goods. When the income innovation is very persistent the sign of the substitution between durable and nondurable goods can be reversed. In particular, a permanent income shock will lead an impatient household to decrease her desired savings because she wants to spend today the whole increase in current income plus a fraction of the expected future income increase. In this situation, the borrowing constraints become more severe, the household substitutes towards nondurable goods, and the response of durable goods to permanent income shocks turns out to be lower than the response of nondurable goods.

The substitution between goods of different durability is consequential for the measurement of consumption insurance in the data. In particular, the observed weak response of nondurable consumption to transitory income shocks cannot be exclusively interpreted as evidence of full insurance because part of the household adjustment is due to the substitution towards durable goods. Likewise, the observed strong response of nondurable consumption to permanent shocks cannot be entirely attributed to lack of insurance because part of it reflects the relative increase of nondurable goods in the consumption basket. Our second contribution is to quantify the different responses of durable and nondurable goods to income shocks of different persistence implied by a standard life cycle model. We use these responses to learn about the amount of insurance available to households of different ages. To do so we calibrate our model to data on wealth accumulation over the life cycle, expenditure share on durable goods, and collateralized borrowing

[4] Luengo-Prado (2006) already showed that the fraction of the durable goods that can be collateralized is quantitatively important in understanding the aggregate consumption response to income shocks.

in durable goods. As over identifying restrictions, the model is quantitatively consistent with the observed responses of nondurable consumption expenditure to transitory and permanent income shocks, as measured by Blundell, Pistaferri, and Preston (2008) in the PSID.

We find that, when taking the average over all the working age population, positive transitory shocks lead to a 9% increase in nondurable consumption and a 22% increase in the stock of durable goods. Given the calibrated share of durable and nondurable consumption expenditure, the whole basket of consumption goods increases by 12%. This implies that the average amount of household insurance to transitory shocks is relatively lower than commonly thought. Furthermore, if we look at the average response among households with heads below age of 40, the model implies a transmission of transitory shocks of 14% for nondurable consumption, of 51% for durable goods, and of 22% for the consumption basket. This leads us to conclude that the amount of insurance against transitory income shocks available to young households is quite low, much less than implied by only looking at nondurable consumption expenditure.

Regarding permanent shocks, the model results imply a transmission of 61% into nondurable goods,a transmission of 53% into durable goods, and a transmission of 59% to the whole consumption basket. This implies that the average amount of household insurance to permanent shocks is slightly higher than commonly thought. If we look at the responses among young households, the model predicts a transmission of permanent shocks of 84% into nondurable consumption, of 60% into durable goods, and of 79% into the consumption basket. This leads us to conclude that the amount of insurance of young households against permanent income shocks is very low, although better than implied by only looking at nondurable consumption expenditure.

One key aspect of our quantitative analysis is the importance of the borrowing constraints. In our main exercise we take a very conservative view in assuming that all household assets are liquid and can be used without any cost to smooth income shocks. However, as argued by Kaplan and Violante (2014), many middle-aged wealthy households may behave as borrowing constrained if the income shocks they suffer are not large enough and their portfolio contains many illiquid assets such as pensions funds or real estate that are costly to adjust. In our robustness section we explore this possibility, and show that considering illiquid assets may increase substantially the amount of substitution between durable and nondurable goods upon arrival of an income change.

The remaining of the paper is organized as follows. We describe the basic model in Section 2. We calibrate the model in Section 3 and present our main quantitative findings in Section 4. Then in Section 5 we discuss several alternative calibrations. Finally, Section

6 concludes.

2 The Model

We use a standard life-cycle model of consumption with idiosyncratic shocks to labor earnings and borrowing constraints as in Huggett (1996), and we extend it to allow for durable as well as nondurable consumption goods.[5] Following Fernández-Villaverde and Krueger (2011), durable goods enter the (homothetic) utility function, they serve as partial collateral for borrowing, and they can be adjusted at no cost.

2.1 Description

Life cycle. Households are born at age $t = 1$ as working adults, retire at age $t = T_R$, and die for sure at age $t = T$. The probability of surviving between age t and $t + 1$ is given by $\pi_{t,t+1}$, and we denote by π_t the unconditional probability of a newborn surviving to age t.

Preferences. Households have time-separable preferences defined over streams of consumption of nondurable goods C_t and service flows of durable goods D_t, which are assumed to be proportional to its stock. The lifetime objective function of a household is given by

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

where $\beta > 0$ is the subjective discount factor. We assume CRRA preferences over a CES aggregator of nondurable consumption and services from durable goods:

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

where $\sigma > 0$ measures the degree of risk aversion, $\epsilon < 1$ measures the elasticity of substitution between goods and $0 < \gamma < 1$ captures the weight of each type of consumption in households' preferences. The stock of durable goods evolve according to the following law of motion:

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

[5] This class of life cycle models has been extensively used to measure the amount of consumption insurance when only nondurable consumption goods are available, see for instance Storesletten, Telmer, and Yaron (2004) or Kaplan and Violante (2010).

where δ is the depreciation rate. 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.

Labor income and pension income We denote the labor earnings at time t by Y_t, and we assume that the stochastic process governing the log of labor earnings, y_t, can be represented as the sum of a random walk z_t with innovations η_t, a purely transitory shock ε_t, and a deterministic age-specific mean μ_t:

$$\begin{aligned} y_t &= \mu_t + z_t + \varepsilon_t \\ z_t &= z_{t-1} + \eta_t \end{aligned} \tag{3}$$

where $\varepsilon_t \sim \mathcal{N}(0, \sigma_\varepsilon)$, $\eta_t \sim \mathcal{N}(0, \sigma_\eta)$, and $z_0 \sim \mathcal{N}(0, \sigma_{z_0})$. After age T_R, households receive an age-invariant payment from the government. This payment is household-specific and it is the sum of an average transfer B, which represents the payments from Medicare, and a function of the household entire labor income history $Y^R(H_{T_R})$, which represents the actual retirement pension from Social Security. At any age, we summarize past earnings history in the variable $H_t = \frac{1}{t}\sum_{j=1}^{t} Y_j$, or recursively:

$$H_t = \frac{(t-1)H_{t-1} + Y_t}{t} \tag{4}$$

Making pension payments a function of the history of past income is important because, as it will be discussed in later sections, the amount of substitution between durable and nondurable goods depends on the persistence of the income shocks. 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 add the Medicare transfer for quantitative reasons. As shown by Huggett (1996), transfers during retirement, if unlinked to labor income, reduce (increase) the incentives to save for poor (rich) households. This is important to help accounting for the observed wealth inequality by age.

Financial markets, borrowing constraints, and budget constraints Households use one-period risk-free bonds to save and possibly borrow at an interest rate r. We denote by A_t the total stock of bonds at the beginning of age t. 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,

$$(1+r)A_{t+1} + \theta\,(1-\delta)\,D_t \geq \underline{A}_t \tag{5}$$

Note that the case with $\underline{A}_t = 0$ and $\theta = 0$ precludes borrowing altogether, whereas the case with $\underline{A}_t = 0$ and $\theta > 0$ allows some collateralized debt. The extreme case of $\underline{A}_t = 0$ 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 Fernández-Villaverde and Krueger (2011). With $\underline{A}_t < 0$ household can access some non-collateralized debt. With all the elements defined we can construct the budget constraint during working life as:

$$C_t + I_t + A_{t+1} \leq (1+r)A_t + Y_t(z_t, \varepsilon_t) \tag{6}$$

and during retirement:

$$C_t + I_t + A_{t+1} \leq (1+r)A_t + Y^R(H_{T_R}) + B \tag{7}$$

2.2 Optimal choices

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 the sequence of budget constraints (6) and (7), the laws of motion defined by (2) and (4), the borrowing constraints (5), the stochastic process for labor income defined in (3), and some initial conditions A_1, D_0, and z_0.

The first order conditions for an optimum are the standard ones,

$$U_C\left(C_t, D_t\right) = \mu_t \tag{8}$$

$$\beta\pi_{t,t+1}(1+r)E_t\left[\mu_{t+1}\right] = \mu_t - \lambda_t \tag{9}$$

$$\frac{U_D\left(C_t, D_t\right)}{U_C\left(C_t, D_t\right)} = \left(\frac{r + \delta + \frac{\lambda_t}{\mu_t}\left(1-\delta\right)\left(1-\theta\right)}{1+r}\right) \tag{10}$$

where λ_t is the the multiplier associated to the borrowing constraint, and μ_t is the multiplier associated to the period budget constraint. Equation (8) equalizes the shadow value of resources within the period to the marginal utility of consumption. Equation (9) is the Euler equation that describes the law of motion of the shadow value of wealth. Equation (10) drives the choice of durable goods. Whenever the borrowing constraint does not bind at t, we have that $\lambda_t = 0$ and equation (10) reduces to the standard condition

$$\frac{U_D\left(C_t, D_t\right)}{U_C\left(C_t, D_t\right)} = \frac{r+\delta}{1+r} \tag{11}$$

This states that the marginal rate of substitution between durable and nondurables is equal to the user cost of durables. Hence, the ratio between marginal utilities is independent of individual level variables and will be equalized across households. 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 the consumption bundle reacts to shocks, but not its composition. Thus, the response of durable and nondurable goods is the same. In fact, under the assumed utility function the consumption ratio is given by

$$\left(\frac{C_t}{D_t}\right)^{1-\epsilon} = \left(\frac{\gamma}{1-\gamma}\right)\left(\frac{r+\delta}{1+r}\right) \tag{12}$$

and this equation can be used in the Euler equation to derive an expression for nondurable 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.[6]

[6] 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 constraint and hence the ratio of multipliers $\frac{\lambda_t}{\mu_t}$. As we argue in the next subsection, a purely transitory positive shock unequivocally alleviates the borrowing constraint, hence λ_t/μ_t falls and there is a substitution towards durable goods and away from nondurable goods. Therefore, the response of durable goods to transitory income shocks is larger than the response of nondurable goods. A permanent shock may, but not necessarily will, have a similar effect. Whether it does or not will depend on the desired path for consumption. In particular, whenever households are impatient (in the sense of desiring a falling consumption profile over time) a positive permanent shock will have the opposite effect, increasing λ_t/μ_t and leading to a substitution towards nondurable goods. In this case, the response of durable goods to the permanent income shock will be smaller than the one of nondurable goods.

2.3 A model without uncertainty

To understand the role of the persistence of income shocks let's simplify our model in a few respects. First, households live forever and survival probabilities are equal to one in all periods; second, there is no retirement and labor earnings Y_t are deterministic and given by the recursion $Y_{t+1} = \rho Y_t$ with $1 > \rho > 0$; and third, there are no borrowing constraints. Under these simplifications the optimal basket of durable and nondurable goods is given by equation (12). Substituting it into the Euler equation we can obtain an expression for nondurable consumption growth,

$$C_{t+1} = [\beta (1 + r)]^{1/\sigma} C_t$$

and for the stock of durable goods,

$$D_{t+1} = [\beta (1 + r)]^{1/\sigma} D_t$$

And substituting this expression in the law of motion for durables we obtain that expenditure on durable goods must grow at the same rate as the stock:

$$I_{t+1} = [\beta (1 + r)]^{1/\sigma} I_t$$

Without liquidity constraints, the relevant resource constraint is given by the intertemporal budget constraint

$$\sum_{j=0}^{\infty} (1+r)^{-j} (C_{t+j} + I_{t+j}) = \sum_{j=0}^{\infty} (1+r)^{-j} Y_{t+j} + (1+r) A_t$$

Assuming $1 + r > \rho$ and $1 + r > [\beta(1+r)]^{1/\sigma}$ for a bounded problem, we can use the expressions for consumption growth above to write total expenditure in t as a function of income and assets in t,

$$C_t + I_t = \frac{1 - \frac{[\beta(1+r)]^{1/\sigma}}{1+r}}{1 - \frac{\rho}{1+r}} Y_t + \left(1 - \frac{[\beta(1+r)]^{1/\sigma}}{1+r}\right)(1+r) A_t$$

Given the expression for expenditure, we can also write A_{t+1} as a function of Y_t and A_t:

$$A_{t+1} = \frac{[\beta(1+r)]^{1/\sigma} - \rho}{1 + r - \rho} Y_t + [\beta(1+r)]^{1/\sigma} A_t \tag{13}$$

These two expressions allow us to understand the effect on savings of an income increase. If $\rho < [\beta(1+r)]^{1/\sigma}$, the marginal propensity to spend out of an increase in Y_t will be positive and less than one, and hence savings in t will increase with Y_t. That is to say, whenever the income growth is less than the desired consumption growth, part of an increase in income today is saved and spread over future periods. Instead, if $\rho > [\beta(1+r)]^{1/\sigma}$ the marginal propensity to spend out of an increase in Y_t is larger than one, and hence an income increase generates a reduction in savings or an increase in borrowing.

Let's now map these results into the full model with life cycle and uncertainty. A purely transitory income shock ($\rho = 0$) increases desired savings, and hence it alleviates the borrowing constraints in case they were binding. By alleviating the borrowing constraints the desired ratio C_t/D_t falls. As we consider more persistent income shocks (larger ρ), the household wants to spend a larger fraction of today's income increase because the higher persistence implies that income will also be larger in the coming periods. Hence, desired savings increase less and there is a smaller reduction in the severity of the borrowing constraint. As a result, the rebalancing between C_t and D_t is also smaller. Finally, whenever the income shocks have a large persistence (large ρ), future income increases almost as much as current income. If the desired consumption growth is less than the persistence of the shock, the household will like to borrow against future income and increase expenditure today by more than the income increase. When borrowing constraints are binding this makes them more severe, hence the ratio C_t/D_t goes up and the response of nondurable goods to the income shock is larger than the response of durable goods.

2.4 The transmission of shocks and the measure of insurance

Let c_{it} be log nondurable consumption for household i at age t and d_{it} be log durable consumption for household i at age t. Following Blundell, Pistaferri, and Preston (2008), we define the transmission coefficients for the shock x_{it} as

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

These coefficients measure the proportional change in each consumption good that arises as a response to shocks. The coefficients ϕ_x^c have been used as a measure of (lack of) insurance because complete markets implies equalization of marginal utilities across states, and in a one-good model this implies equalisation of consumption levels. Then, under complete markets $\phi_x^c = 0$, and the larger ϕ_x^c, the further away from complete markets. With two goods, however, ϕ_x^c might be very small while ϕ_x^d is large, or the other way around. This will happen whenever income shocks give rise to substitution between goods. To come up with a measure of lack of insurance that can be used in the two-good model as well as in the one-good model we consider the transmission of income shocks into the consumption basket $V = C^\gamma D^{1-\gamma}$,

$$\phi_x^v = \frac{cov\left(\Delta v_{it}, x_{it}\right)}{var\left(x_{it}\right)} = \phi_x^c + (1-\gamma)\left(\phi_x^d - \phi_x^c\right) \tag{14}$$

where v_{it} is the logarithm of V_{it}. Equation (12) above shows that in the absence of binding borrowing constraints $c_{it} - d_{it}$ is independent from shocks. Hence, $\phi_x^c - \phi_x^d = 0$ for both shocks and $\phi_x^v = \phi_x^c$. In that case, the transmission coefficient of nondurable consumption is a correct measure of lack of insurance. Instead, when the borrowing constraints bind, equation (10) shows that $\phi_x^d - \phi_x^c > 0$ if the shock x_{it} alleviates the borrowing constraint and $\phi_x^d - \phi_x^c < 0$ if the shock makes the borrowing constraint more severe. In this situation, ϕ_x^c gives a biased measure of insurance and the difference $\phi_x^c - \phi_x^v$ tells us how much of the transmission of income shocks into nondurable consumption is due to substitution between goods. In particular, note that this bias is given by

$$\phi_x^c - \phi_x^v = (1-\gamma)\left(\phi_x^c - \phi_x^d\right)$$

Therefore, ϕ_x^c is a better measure of insurance (i.e. the bias is smaller) whenever the substitution between goods $\left(\phi_x^c - \phi_x^d\right)$ is smaller and/or the expenditure share of nondurable goods γ is bigger.

3 Calibration

We need to set the values for 11 parameters. The key parameters β, γ, and θ are calibrated such that the model is consistent with the data on wealth accumulation over the life cycle, the expenditure share on durable goods, and the amount of borrowing collateralized with durable goods. The rest of parameters are set directly from common values in the literature. As we will see, the model is also quantitatively consistent with an important fact that we do not target. In particular, it delivers the observed average transmissions of transitory and permanent shocks into nondurable consumption as measured by Blundell, Pistaferri, and Preston (2008). This property is central to the question at hand because it implies that the model is consistent with the average amount of insurance to transitory and permanent shocks as measured in the data.

3.1 Data

We use the Panel Study of Income Dynamics (PSID) to measure the distribution of wealth holdings and the Consumer Expenditure Survey (CEX) to measure the aggregate composition of the consumption basket and the extent of borrowing against durable goods. Because we want to compare the responses to shocks of nondurable consumption to the ones measured by Blundell, Pistaferri, and Preston (2008) in the data, the reference period is 1980-1992. In this time interval, we have only two PSID waves with data on wealth: 1984 and 1989. We also use the 1994 wave, which is close enough in time. In our PSID sample, we include only married households with a head active in the labor market. Our measure of wealth in the PSID data is given by total net worth: this is the value of all assets, including housing, minus all debts. Net worth includes the value of vehicles (minus the outstanding debt associated to their purchase), but not the value of other durable goods. We will associate wealth in the data to the sum $W = A + D$ in the model.[7]

As for the CEX, we work with the series of annual cross-sections described in Harris and Sabelhaus (2000). We classify the different expenditure categories in the CEX as either durable or nondurable. Durable goods include cloth, jewellery, furniture, household

[7]Hence, in our main calibration houses are going to be part of financial wealth A. However, in Section 5 we explore a calibration in which houses are part of the stock of durable goods that yield utility.

appliances, vehicles and spare parts, books, and sport and recreational equipment, but exclude housing. Nondurables include food and other household supplies, household utilities, services, public transport fees, fuel and tolls expenditures.[8] 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 therapeutic equipment from the total stock of consumer durables.

3.2 Parameter values

Table 1: Calibration targets and results

Parameter	Value	Target/Source	Model	Data
Common				
σ	2			
r	3%			
Income process				
σ^2_ε	0.05			
σ^2_η	0.01	Kaplan-Violante (2010)		
$\sigma^2_{z_0}$	0.15			
Retirement Income				
Initial wealth				
Shape parameter	1.1539	Gini of W at 21-25 (PSID)	0.8468	0.8469
average A_1	1.077	average W at 21-25 (PSID)	1.077	1.077
With durables				
β	0.9854	average W at 56-60 (PSID)	32.875	32.874
γ	0.7870	average I/C (CEX)	0.2444	0.2444
δ	0.1292	aggregate I/D (BEA)	0.1445	0.1400
θ	0.1183	average $\theta\frac{(1-\delta)}{1+r}$ (CEX)	0.1000	0.1000
Without durables				
β	0.9896	average W at 56-60 (PSID)	32.873	32.874

Timing and demographics. A period is a year. We assume households are born to working life at age 25 and retire at age 65. Certain death takes place at age 95. This implies $T_R = 40$ 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.

[8] 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. We exclude health and education expenditures from the analysis as they can be seen more like an investment.

Initial wealth distribution. Households enter the model with some random initial wealth A_1, which is drawn from an exogenous distribution. We use a Pareto distribution to capture the skewness of the observed distribution for wealth at young ages in our PSID sample. The parameters of the distribution function are chosen to replicate the average wealth and the degree of inequality, as measured by the Gini index, for households in the 21-25 age bracket.

Income process. We calibrate the earnings process with the values used by Kaplan and Violante (2010). These authors use a PSID sample as in Blundell, Pistaferri, and Preston (2008). They calibrate the deterministic component to mimic the average age profile of after-tax earnings, and the the variance of the permanent shock σ_η^2 to match the increase in residual earnings dispersion over the life cycle. The variance of the transitory shocks σ_ε^2 (set to 5 times the one of the permanent shocks) is taken directly from Blundell, Pistaferri, and Preston (2008) estimates. The initial variance of the permanent component of income $\sigma_{z_0}^2$ is then set to replicate the variance of dispersion of residual earnings at age 25.

Retirement income. The pension benefit is a concave function of average working life earnings, explicitly capturing the progressivity of the U.S. social security system. We parametrize it following Storesletten, Telmer, and Yaron (1999).[9] The Medicare payment B is calibrated to the ratio of total Medicaid payments to individuals aged 65+ and the number of households at that age range. This gives a payment of \$8,641 per household and year.[10]

Preferences. Our utility function has three parameters to be set: ϵ, which captures the elasticity of substitution between goods; σ, which measures the coefficient of relative risk aversion; and γ, which measures the weight of nondurable goods. In addition, we have the intertemporal discount factor β. We set σ to 2, as widely used in the literature. We also fix $\epsilon = 0$, implying a Cobb-Douglas aggregator for durables and nondurables. Both

[9]This function is characterized by a minimum and a maximum level of benefits, and a piecewise-linear function of average earnings in between. 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.

[10]According to MCBS Project (2006), in 2002 total Medicare payments to beneficiaries 65+ amounted to \$229,915 million (Table 4.1), while the number of beneficiaries aged 65+ were 35,954,880 individuals (Table 1.1). Of those, 52% were married (Table 2.1). Assuming that the 65+ married individuals are married to other 65+ married individuals, this gives 26,606,611 households aged 65+.

aggregate and micro data are consistent with this choice.[11] We choose γ to match the average durables to nondurables expenditure ratio in the CEX. The discount factor β is set by matching the average wealth at ages 56-60 in the PSID. Since we feed the wealth levels at young ages into the model, this target captures the average increase in wealth over the life cycle.

Technology parameters. The return to savings r is set to 3%. The depreciation rate δ is set to 0.13. We follow Aaronson, Agarwal, and French (2012) in using this value computed by Campbell and Hercowitz (2003) for consumer durables excluding housing. We consider alternative values in section 5.2.

Financial markets parameters. A key parameter in our analysis is θ, which captures the extent to which durable goods can be used as collateral for borrowing. We exploit the information available in the CEX on new loans acquired to purchase vehicles. In particular, for each household with positive expenditure on durable goods, we divide the amount borrowed to purchase vehicles by total expenditure on 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 0.1. In the model, the maximum possible debt to acquire one unit of durable good is given by $\theta\frac{(1-\delta)}{1+r}$, so for a given depreciation rate and a given interest rate, we choose the θ that makes this expression equal to 0.1. In our baseline calibration we set $\underline{A}_t = 0$, which precludes unsecured debt. This choice corresponds quite closely to the natural borrowing limit in the model given that we do not impose any lower bound on the log of the income process.[12] In the robustness section we check for the effects of $\underline{A}_t < 0$.

Model without durables. We also calibrate a version of the model without durable goods. In this case θ, δ, and γ are absent, and we recalibrate β to keep the wealth profiles

[11] The aggregate time series data from the US shows that, while the relative price of durable goods to nondurable goods has fallen steadily for the last 40 years, the share of durables to nondurables has remained stable over time. This feature is consistent with a unit elasticity of substitution between goods. The attempts to estimate the elasticity of substitution between durables and nondurables using micro data also support the specification that the time series evolution of aggregate data suggests. Most studies cannot reject the hypothesis of $\epsilon = 0$. See Fernández-Villaverde and Krueger (2011) for a summary of empirical estimates of this parameter.

[12] Indeed, given that labor earnings can be zero, the natural borrowing restricts the borrowing limit to be equal to the present discounted value of pension benefits, which has a positive lower bound given by the existence of a minimum benefit. We find this channel for unsecured debt as an artefact of the model, rather than as a meaningful economic mechanism, and hence choose to switch it off. We discuss the implications of enabling it in Section 5.3.

as in data.

3.3 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. Red lines depict the profiles emerging from the model without durable goods, while the blue lines represent our baseline model with durables.

Figure 1: Average Life Cycle Profiles

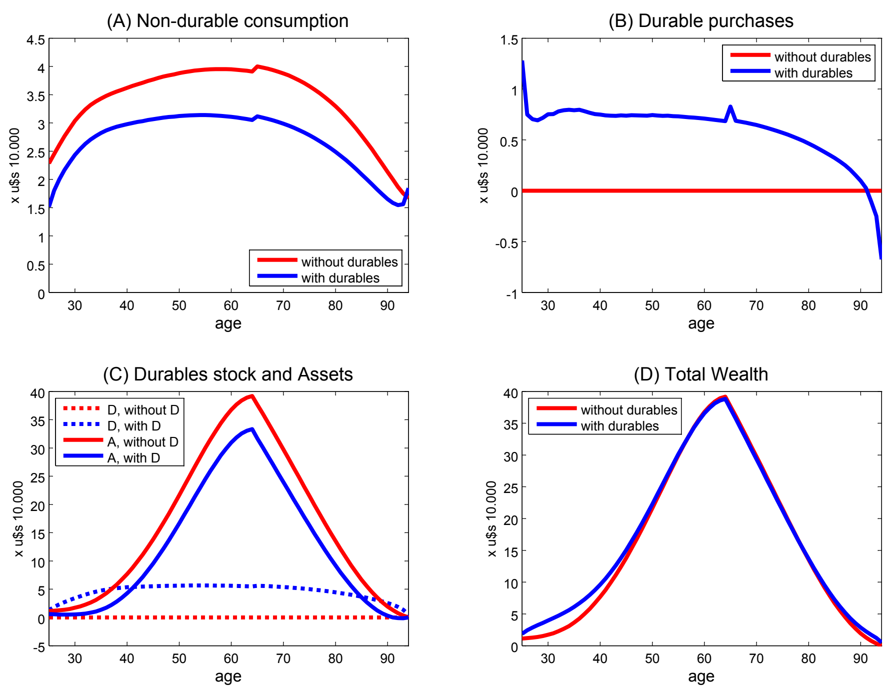

Nondurable consumption (top left panel) is hump-shaped as in the data. It peaks around 55 years of age, somewhat later than documented by Gourinchas and Parker (2002) (45 years old) and Fernández-Villaverde and Krueger (2006) (50 years old). The size of the hump in nondurable consumption exceeds its empirical counterpart: nondurable consumption roughly doubles between age 25 and the peak, compared to the estimated increase of 50% in Gourinchas and Parker (2002) or the 25%-40% increase in Fernández-Villaverde and Krueger (2006). A similar hump-shaped pattern is observed for expenditure on durable goods (top right panel), apart from an initial spike in durables expenditure because simulated households are born without any durables. Fernández-Villaverde and Krueger (2006) document a hump-shaped profile for durables expenditure, with a similar timing of the hump and the peak being approximately 33% higher than the minimum, in contrast with the 15% increase obtained by our model. Total expenditure (not reported) combines the excessive hump of nondurables and the moderate hump of durables: the increase in total expenditure over the life cycle is about 40% in the model, compared to the 30% estimated by Gourinchas and Parker (2002) and Fernández-Villaverde and Krueger (2006).[13]

Wealth accumulation (bottom right panel) follows the characteristic pattern in this family of models, with most of the asset accumulation taking place near retirement. The size of the peak is matched to the data by calibration.[14] This is important. Models that are consistent with the size of the life cycle hump in consumption largely understate the life cycle hump in wealth. For instance, in Kaplan and Violante (2010) the hump in wealth is half as large. Finally, the wealth composition in the economy with durables (bottom left panel) shows how households build up a stock of durables at early stages of the life cycle (dotted blue line) at the expense of accumulating financial assets (solid blue line). This feature is analogous to the findings in Fernández-Villaverde and Krueger (2011).

Figure 2: Average of Euler Equations

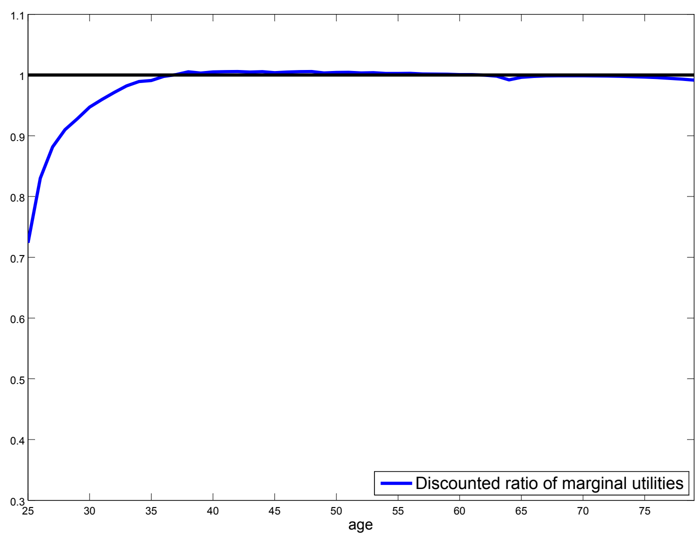

Figure 2 illustrates the incidence of borrowing constraints over the life cycle. Specifically, it depicts the cross-sectional average by age of the ex-post discounted ratio of marginal utilities, $\beta\pi_{t,t+1}(1+r)\frac{\mu_{t+1}}{\mu_t}$. This is just an ex-post version of the Euler equation,

[13] Figure 1 shows an additional spike in consumption of both goods at retirement. During retirement there is no uncertainty left in the model, so precautionary motives for saving disappear after T_R.

[14] The quick deaccumulation of assets during retirement is somewhat counterfactual, see for instance Nakajima and Telyukova (2012). Since we focus on the transmission of income shocks during working life, we abstract from motives to save during retirement, such as health uncertainty or intentional bequests.

and should be equal to 1 in the absence of binding borrowing constraints. We can see that the largest deviations from the optimal intertemporal allocation of consumption are concentrated among young households. By the age of 40, the borrowing constraint is not binding for most households, and, on average, no significant deviation from the desired allocations are observed.

Figure 3: Wealth inequality over the life cycle

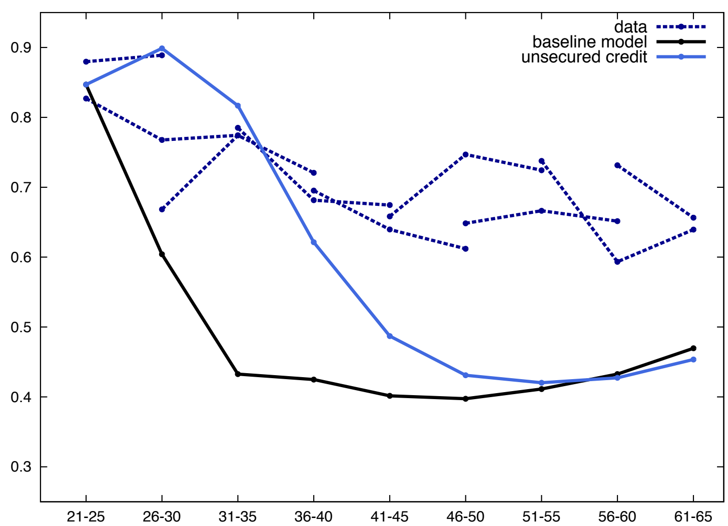

Note: Data corresponds to the PSID waves of 1984, 1989, and 1994. Each segment corresponds to the average wealth inequality for a given cohort, followed over the three years.

Finally, Figure 3 shows the evolution of wealth inequality. The model is able to capture the decline in inequality over the life cycle, albeit at a higher pace than in the data. The difference in inequality levels between the model and the data is attributable to the absence of households with little or no wealth. In the absence of unsecured debt, the model is unable to generate households with no wealth because all households hold a positive stock of durable goods, and they can borrow only against a fraction of this stock. In Section 5.3 we study an economy with unsecured borrowing, which delivers wealth inequality statistics closer to the data.

4 Results

We use our calibrated model to compute the transmission coefficients with simulated data. Table 2 summarizes the results. Panel (a) reports the average transmission coefficients for all households, panel (b) shows the same information for the youngest households only, and panel (c) for the oldest (working-age) households.

Table 2: Average transmission coefficients of income shocks to consumption (%)

	ϕ_x^c	ϕ_x^d	$\phi_x^c - \phi_x^d$	ϕ_x^v	$\phi_x^c - \phi_x^v$
(a) All households					
Transitory shocks	8.6	21.9	-13.3	11.5	-2.8
Permanent shock	60.9	52.5	8.4	59.1	1.8
(b) Young households (below 40)					
Transitory shocks	14.2	50.7	-36.4	22.0	-7.8
Permanent shock	83.5	60.4	23.2	78.6	4.9
(c) Old households (over 40)					
Transitory shocks	5.5	5.8	-0.3	5.5	-0.0
Permanent shock	48.3	48.2	0.1	48.2	0.0

Note: The first column reports the transmission of income shocks into nondurable goods; the second column reports the transmission into durable goods; the third column reports the substitution between durable and nondurable goods as a response to the income shocks; the fourth column reports the transmission of shocks into the consumption basket and hence it reflects our measure of lack of insurance; the fifth column reports the bias in the standard measure of consumption insurance.

For transitory shocks, we find that 8.6% of innovations are transmitted into nondurable consumption. In comparison to the empirical findings by Blundell, Pistaferri, and Preston (2008), our model generates slightly larger transmission of transitory shocks to nondurable consumption, roughly 3 percentage points higher than their baseline estimate of 5.3%, but this is within one standard deviation of their estimate. As we discuss in more detail below, the transmission of shocks into nondurable consumption expenditure is larger for the young, averaging 14.2% among households below age 40. Regarding permanent shocks, our model generates about the same transmission to nondurable consumption as found in the data (60.9% compared to 64.2% in Blundell, Pistaferri, and Preston (2008), with a standard deviation of 9%).[15] The transmission of the permanent shocks is also much higher among the young (83.5%) than among the old (48.3%)

Our main finding is that the transmission of income shocks into durable goods is substantially different from the transmission into nondurables, and that these differences go in opposite direction depending on the type of shock. Regarding transitory shocks, the response of durables for the overall population is 21.9%, which is 13.3 percentage

[15]The estimation method used by Blundell, Pistaferri, and Preston (2008) is known to bias upwards the transmission of permanent shocks in the presence of binding borrowing constraints, see Kaplan and Violante (2010). We applied their estimation procedure to our simulated data and obtained a point estimate of 72.7, still within the range of their empirical estimates.

Figure 4: Age profiles of transmission coefficients of income shocks to consumption

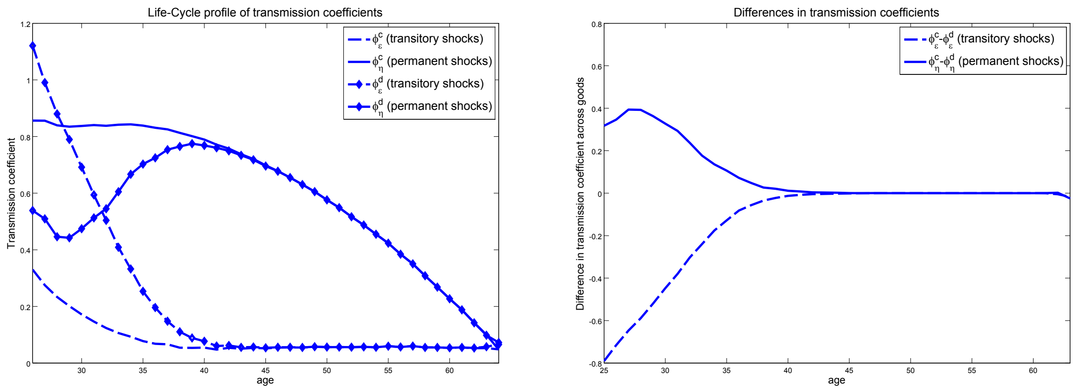

points larger than the response of nondurables. This reflects a shift of the consumption basket towards durable goods due to the loosening of borrowing constraints. Instead, the response of durables to permanent shocks is 52.5%, which is 8.4 points smaller than the one of nondurables. This reflects a shift away from durable goods as households increase their desire to borrow from the higher expected future income. The substitution of consumption goods induced by the shocks changes over the life cycle along with the incidence of the borrowing constraints. In particular, for young households the transmission of transitory shocks into durable goods is 36.4 percentage points higher than into nondurables, and the transmission of permanent shocks into durable goods is 23.1 percentage points lower than into nondurables.

These differences in transmission between goods imply non-negligible biases in the insurance measures based on nondurable consumption, especially for the young. We find that for the overall population the transmission of transitory shocks into the consumption basket, ϕ^v_ε, is 11.5% (see column 4 in Table 2), which is 2.8% points larger than the transmission into nondurable consumption, ϕ^c_ε (see column 5). For the young, the transmission into the consumption basket is 22.0%, which is 7.8% points larger than the transmission into nondurable consumption. Therefore, we conclude that transitory income shocks are not easy to insure for the young, despite the fact that by looking at the response of nondurable consumption it seems so. Regarding the permanent shocks, the bias of the transmission of permanent shocks is 1.8% for the overall population and 4.9% for the young.

For a more detailed assessment of the role of age in consumption insurance, in the left panel of Figure 4 we plot the transmission coefficients by age for both nondurable, ϕ^c_x (plain lines), and durable goods, ϕ^d_x (lines with diamonds). The solid lines correspond to the permanent shock and the dashed lines to the transitory shock. As shown by Kaplan and Violante (2010), this type of models predict a clear life-cycle pattern in the transmission of shocks. For both types of income shocks, the transmission to nondurable expenditure 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.[16] The shape of the age profile for transmission is the result of two forces: the age profile of binding borrowing constraints and the proximity 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.

The difference between the transmission to nondurables and durables, $\phi^c_x - \phi^d_x$, gives our measure of substitution between goods for each type of shock, and is plotted in the right Panel of Figure 4. We see that the substitution between consumption goods 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 substitution towards durable goods ($\phi^d_\varepsilon > \phi^c_\varepsilon$), hence the response of nondurables ϕ^c_ε is lower than the response of the consumption basket ϕ^v_ε. A permanent shock to a constrained household has the opposite effect, substituting consumption towards nondurable goods ($\phi^d_\eta < \phi^c_\eta$), and hence nondurable goods react more than the composite basket. In terms of equation (12), this implies that λ_t/μ_t, our measure of the tightness of the borrowing constraint, co-moves 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.

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

4.1 Differences between models

In this Section we want to compare the transmission coefficients of nondurable consumption predicted by our model to those of the model without durable goods. This transmissions may be quite different because the two models differ not only in the substitution between goods but also in some other important aspects. First, households in the model with durable goods have a higher capacity to borrow (as long as $\theta \neq 0$). Second, while the two economies are calibrated to the same total wealth, the timing of wealth accumulation is different: in the model with durable goods households accumulate more wealth in the first part of the life-cycle, whereas in the model without durable goods households have more wealth in the second part of the life cycle (see Figure 1). Third, the composition of wealth is also different: in the model with durable goods, most of the assets held before the age of 40 are durable goods. Finally, in order to achieve the same amount of total wealth in both economies, households are somewhat more impatient in the model with durables (see Table 1).

Figure 5: Transmission coefficients (nondurables)

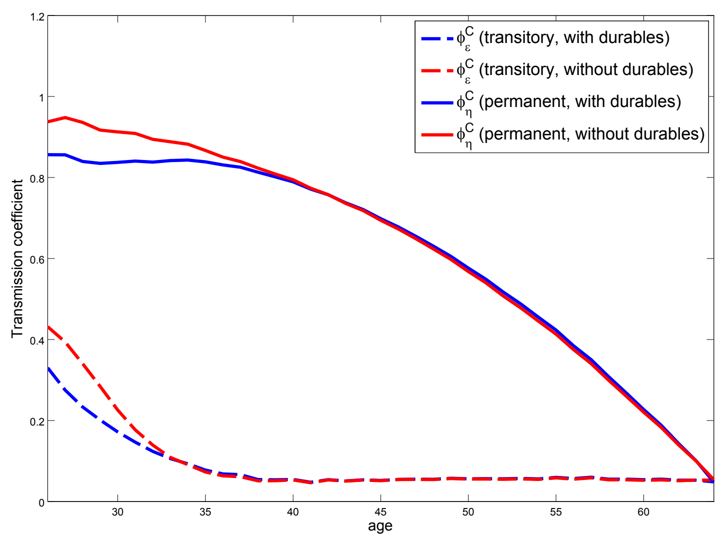

In Figure 5 we plot the transmission coefficients for nondurable consumption for both models. Blue lines depict the transmission coefficients in the model with durable goods, while red lines indicate the transmission coefficients in the model without durables. The solid lines correspond to the transmission of permanent shocks, and the dashed lines to the transitory shock. We see that the main differences between models are in younger households. In particular, the transmission of both shocks to nondurable consumption is lower in the model with durables up to age 40, being practically identical thereafter. The transmission of the transitory shock is lower in the model with durables because of the substitution towards durable goods. The transmission of the permanent shock is lower in the model with durables despite the fact that there is substitution towards nondurable goods. However, in the model with durables households are better self-insured and a lower fraction of the shocks is transmitted into consumption.

5 Extensions and robustness checks

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.

Table 3: Alternative parameterizations: calibration results.

Specification	Baseline	Illiquid wealth	Housing	Unsecured debt
Parameters				
β	0.9854	0.9666	0.9848	0.9840
γ	0.7870	0.7780	0.6650	0.7828
θ	0.1183	0.1183	0.8000	0.1183
δ	0.1292	0.1292	0.0452	0.1292
τ	-	0.0976	-	-
α	-	-	-	0.4905
Statistics				
average W at 56-60	32.875	32.874	32.875	32.872
fraction with $W < \$1000$	0.0000	0.0000	0.0000	0.1201
average I/C	0.2444	0.2444	0.4510	0.2444
aggregate I/D	0.1445	0.1441	0.0619	0.1456
average W^{liq}/W^{noliq}	-	0.2279	-	-

Note: The first column corresponds to the baseline economy; the second, third, and fourth columns correspond to the extensions analysed in sections 5.1, 5.2, and 5.3 respectively.

5.1 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 in terms of liquidity. According to Kaplan and Violante (2014), a fraction as large as 80% of average household wealth is held in illiquid assets, which means that a large fixed cost has to be paid in order to use it. Therefore, the self-insurance role of wealth may be overstated in our baseline exercise. Modelling portfolio decisions in the presence of assets that differ in liquidity and rate of return 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, 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 retirement 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. 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.

Figure 6: Average Life Cycle Profiles (Illiquid Wealth)

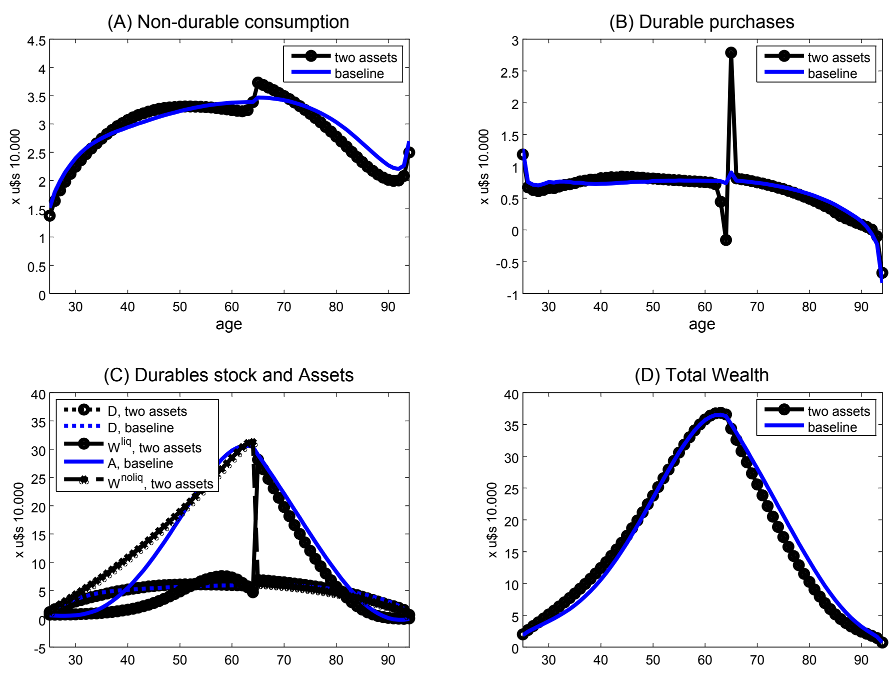

In terms of calibration, introducing illiquid wealth requires to pin down the composition of wealth, which is governed by τ. We hence calibrate β, γ, and τ jointly to match the same statistics as in the baseline case plus the liquid to illiquid wealth ratio for working age households, which is 0.23 as reported by Kaplan and Violante (2014). To be consistent with the definition used in the data, we compute net liquid (illiquid) wealth in the model as total liquid (illiquid) assets minus total liabilities associated to the purchase of liquid (illiquid) assets. In particular, we define liquid wealth W_t^{liq}, and illiquid wealth W_t^{noliq}, as:[17]

$$
\begin{aligned}
W_t^{liq} &= A_{t+1} + \theta \frac{1-\delta}{1+r} D_t \\
W_t^{noliq} &= (1 - \theta \frac{1-\delta}{1+r}) D_t + \underbrace{\tau \sum_{j=1}^{t} (1+r)^{t-j} Y_j}_{\text{retirement accounts}}
\end{aligned}
$$

We obtain a value of τ equal to 9.76%. Figure 6 shows that the age profiles of consumption and wealth in the model with illiquid assets are similar to the ones in the benchmark economy (except for a spike in consumption expenditures at retirement that results from households cashing in their retirement accounts). The main difference is that the economy with illiquid assets features 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.[18]

We report the transmission coefficients by age in Figure 7. We see that, while early in life the amount of rebalancing in the two-asset economy is similar to the one-asset economy, it fades away much later in life in the two-asset case, leading to a higher overall incidence of rebalancing and a flatter age profile for transmission coefficients. This reflects the presence of rich constrained households in the two-asset model. These are households in the second half of their working life, who own a significant amount of assets, but are nevertheless constrained in terms of liquidity, since most of their wealth is accounted by the illiquid asset. The average transmission coefficients generated by the two-asset model are reported in the second 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 baseline model, the transmission of transitory shocks into nondurable consumption increases from 8.6% to 14.4%, whereas the transmission of

[17]Since W_t^{noliq} captures the net value of illiquid assets, it is unclear how the borrowing limit should be specified in the two-asset case. We choose to maintain the restriction on liquid assets defined by (5), so the results are directly comparable to our main exercise.

[18]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 (2014) for a recent analysis of wealth composition when all assets are chosen endogenously.

Figure 7: Difference in Transmission Coefficients (Illiquid Wealth)

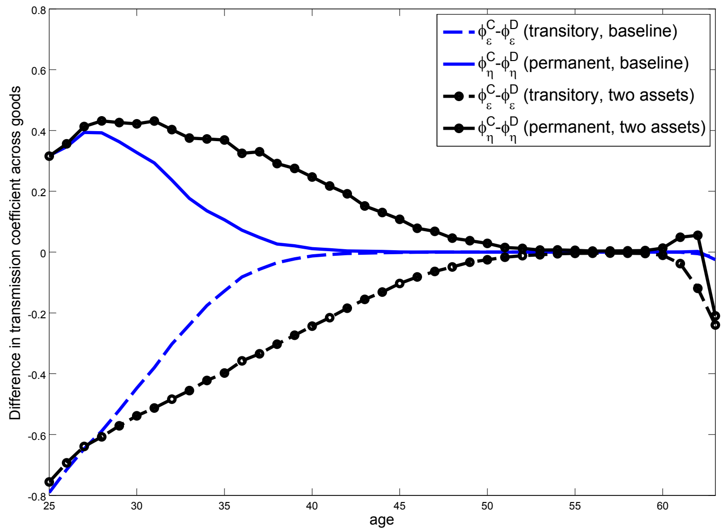

permanent shocks increases from 60.9% to 71.8%. A closer examination of panels (b) and (c) of Table 4 reveals that most of the differences come from older households, especially for permanent shocks: while the differences in transmission coefficients to nondurable goods are hardly above 6 percentage points for households below 40 years old, they can be as large as 17 percentage points for older households. In terms of the transmission of shocks to the consumption bundle, however, there are significant differences throughout the life cycle, since the response of durable goods is substantially different in the two-asset model even at young ages. Therefore, both age groups contribute to the increase in the overall bias relative to the benchmark.

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. From an empirical point of view, our preliminary exercise highlights the importance of distinguishing constrained households in terms of access to liquidity, rather than in terms of levels of net worth.

5.2 Housing

Our main exercise is focused on consumer durables such as cars, furniture, appliances, and smaller durable goods. Our simple framework captures, in a stylized way, the main features of these goods, but it may be a poorer approximation to the characteristics of

Table 4: Alternative parameterizations: transmission of income shocks.

Specification	Baseline	Illiquid wealth	Housing	Unsecured debt
(a) Transmission (all)				
ϕ_ε^c	8.6	14.4	8.9	6.9
ϕ_η^c	60.9	71.8	61.0	62.4
ϕ_ε^d	21.9	37.7	18.0	13.9
ϕ_η^d	52.5	54.4	55.4	56.2
$\phi_\varepsilon^c - \phi_\varepsilon^v$	-2.8	-5.2	-3.1	-1.5
$\phi_\eta^c - \phi_\eta^v$	1.8	3.9	1.9	1.4
(b) Transmission (young)				
ϕ_ε^c	14.2	21.3	14.8	9.2
ϕ_η^c	83.5	83.6	83.3	85.7
ϕ_ε^d	50.7	71.7	39.8	28.5
ϕ_η^d	60.4	46.0	67.7	68.5
$\phi_\varepsilon^c - \phi_\varepsilon^v$	-7.8	-11.2	-8.4	-4.2
$\phi_\eta^c - \phi_\eta^v$	4.9	8.4	2.3	5.2
(c) Transmission (old)				
ϕ_ε^c	5.5	10.5	5.5	5.6
ϕ_η^c	48.3	65.3	48.6	49.4
ϕ_ε^d	5.8	18.6	5.7	5.7
ϕ_η^d	48.2	59.2	48.6	49.4
$\phi_\varepsilon^c - \phi_\varepsilon^v$	-0.0	-1.8	-0.0	-0.0
$\phi_\eta^c - \phi_\eta^v$	0.0	1.4	0.0	0.0

Note: Transmission coefficients are expressed in percentage points. The first column corresponds to the baseline economy; the second, third, and fourth columns correspond to the extensions analysed in sections 5.1, 5.2, and 5.3 respectively.

housing. While the housing stock is an unlikely margin of adjustment to income shocks, for the sake of completeness we perform an alternative calibration including housing in the bundle of durable goods. In practice, this amounts to a reassessment of the durability of goods, δ, the share of expenditures devoted to durables, $1-\gamma$, and the required down payment on durables, θ. It also changes the definition of assets A, as houses now go into D, but it does not change the total net worth $W = A + D$.

Equation (10) states that a lower depreciation rate increases the substitution between goods for a given change in the severity of the borrowing constraints. In our baseline calibration, following Campbell and Hercowitz (2003) we used a value for δ of 12.92%, which is consistent with the aggregate ratio between expenditures and stocks of consumer durables excluding housing. When housing is included, these authors report $\delta = 4.52\%$, which we pick for the present exercise.

Figure 8: Average Life Cycle Profiles (Housing)

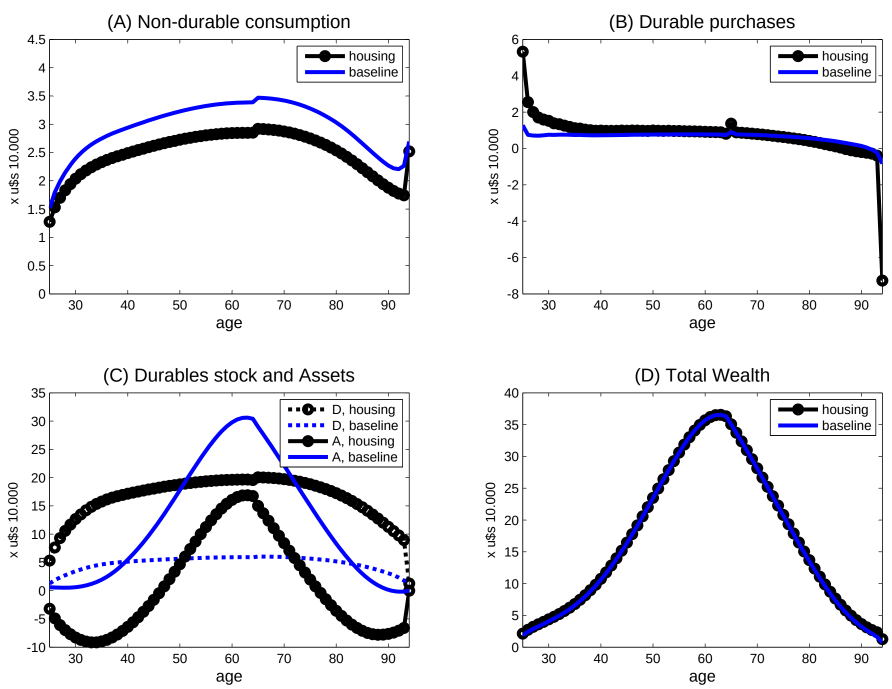

At the same time, equation (10) shows that substitution between goods is less important the lower is the down payment $(1-\theta)$. Our benchmark calibration featured high down payments, as θ was slightly lower than 0.12. This was a result of bundling goods with a high collateral value, such as cars, with many other goods with practically no collateral value. When we include housing, a highly collateralizable asset that accounts for a large fraction of durable goods holdings, the ability of households to borrow against durables has to increase significantly. Following the literature on housing, we set a down payment of 20%, which amounts to $\theta = 0.80$.[19]

Finally, equation (10) implicitly links substitution between goods to the expenditure share of durable goods, $1 - \gamma$. Moreover, equation (14) shows that, for a given degree of rebalancing, higher expenditure shares in durables imply higher differences between the response of nondurable goods to shocks and the response of the consumption bundle. We augment our measure of expenditures on durables in the CEX to include house purchases as well as expenditures in house reforms, leading to a I/C ratio of 0.45. We use this value as a target to calibrate γ in equilibrium.

We recalibrate our model economy accordingly, keeping the rest of the targets at their baseline level. The results of this calibration (summarized in the second column of Table 3) show a minor change in the discount factor β, while exhibiting a substantial reduction in γ. Figure 8 shows no sign of significant changes in the timing of wealth accumulation. The increase in γ and θ and the decrease in δ lead to a higher demand for durables, which translates into a shift in the composition of wealth, reinforcing the additional debt capacity implied by the lower down payment requirement.

Figure 9: Difference in Transmission Coefficients (Housing)

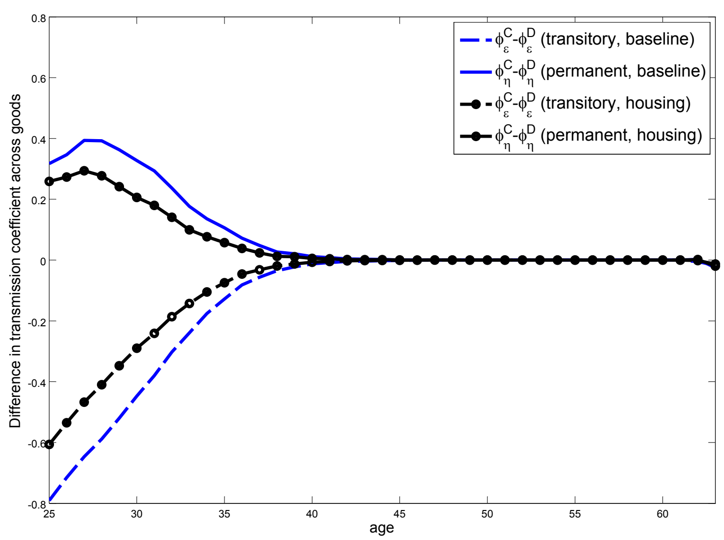

Figure 9 illustrates the amount of substitution between goods, $\phi_x^c - \phi_x^d$. There is a

[19]For instance, this is the value used by Díaz and Luengo-Prado (2010) and Campbell and Cocco (2003). Cocco (2004) argues that reasonable values for housing down payments are between 0.10 and 0.20, while Yang (2009) explores values between 0.10 and 0.50.

sizeable reduction in the extent of the rebalancing throughout the first half of the life cycle. These results are due to the large increase in θ, which more than compensates the lower substitution between goods induced by a lower depreciation rate δ and a higher taste for durables γ. Remarkably, there are only minor differences in the transmission coefficients for nondurables. A comparison of the first and the third columns of Table 4 reveals that the main difference between the baseline economy and the economy with housing is the level of transmission of shocks to durable goods. Importantly, the reduction in rebalancing caused by a higher θ is roughly compensated by the reduction in γ, resulting in ϕ^c being a similarly biased measure of insurance in both economies.

The main conclusion of this exercise is that the response of nondurables to shocks does not depend much on whether housing is included in the definition of durable goods, but its accuracy as a measure of overall insurance does. At the same time, we are cautious about concluding that housing is not important in order to study the transmission of income shocks to consumption. As the results in the previous section suggest, a careful consideration of the liquidity of housing as an asset may lead to a largely different result.

5.3 Unsecured borrowing

Our baseline model requires all debts to be backed by collateral, with θ measuring the tightness of the collateral constraint. However, it precluded unsecured borrowing altogether. This can be rationalized as the equilibrium outcome from a limited enforceability problem where the seizure of tangible assets is the only punishment for defaulting households, with $(1-\theta)$ measuring the loss given default. In this section, we relax this assumption by allowing households to borrow up to a fraction α of their natural borrowing limit, defined as the maximum amount of debt that can be paid back with probability 1. Formally, we define $\underline{A}$ as

$$\underline{A}_t = -\alpha \left[\underline{Y}_{t+1} + \sum_{j=t+2}^{T} \left(\prod_{k=t+2}^{j} \frac{1}{1+r_k} \right) \underline{Y}_j \right] \tag{15}$$

where $\underline{Y}_j$ denotes the lowest possible income for age j given the information available at age t. The baseline model can be regarded as a particular case arising when $\alpha = 0$. When allowing for unsecured debt, we model A_t as an annuity yielding the age-specific return r_t. This return is the result of adjusting the interest rate by the survival probability at each age t. We do this to keep consistency between the assumption that life spans are stochastic, and the assumption that lenders recover $1 + r$ times the amount lent with certainty.[20]

In order to calibrate α, we focus on the fraction of the population with zero or negative net worth in the US.[21] We acknowledge, however, that the level of wealth held by a household is far from being a perfect measure of liquidity constraints.[22]

Figure 10: Average Life Cycle Profiles (Unsecured Debt)

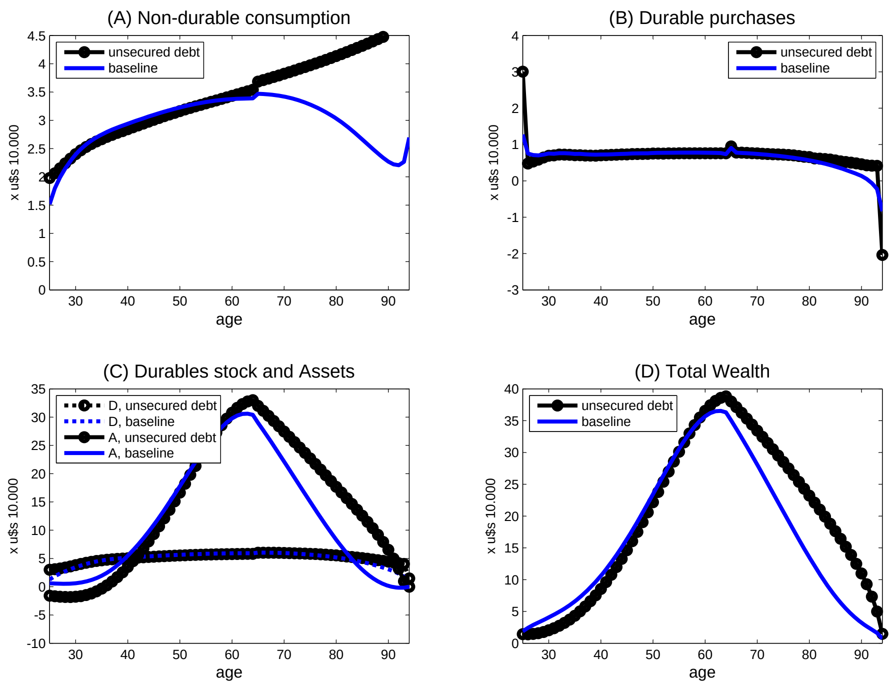

The fourth column of Table 3 shows the results of this calibration. The taste for nondurable goods is unaffected by the change in α. The looser borrowing limit induces lower precautionary savings with respect to our benchmark, decreasing the incentives for young households to accumulate financial assets. However, the introduction of annuities increases the incentives to save at old ages. These two forces combined lead to a discount factor very close to that of the baseline calibration.

Figure 10 shows the role of unsecured debt over the life cycle. Nondurable consumption

[20]Notice that, given our specification for the income process, labor income can be arbitrarily close to zero, while pension benefits are bounded by a minimum level of benefits and Medicare transfers. However, as a household grows older, its earning history is built into its future pension benefits, effectively increasing the natural borrowing limit. Hence, unsecured debt limits increase during working life, more so for luckier households in terms of labor income realizations.

[21]In practice, we set a $1,000 threshold to separate cash holdings from actual savings.

[22]As discussed in Section 5.1, owning assets does not necessarily grant having access to liquidity whenever needed.

grows slower in the economy with unsecured borrowing compared to the baseline economy. Moreover, it is strictly increasing in age, rather than hump shaped. This is a consequence of both the additional borrowing capacity and the incentives to save through annuities. Using unsecured debt, young households can consume more from the beginning of their lives, reducing the growth rate of consumption. After retirement, the return to annuities increases with age, compensating for the decline in survival probabilities. This delays dissaving and leads to an upward-sloping consumption path.

At the same time, allowing for unsecured debt exacerbates the counter-factual initial spike in expenditure on durables, enabling a faster accumulation of durables at the initial stages of working life. Overall, total expenditure is higher early in life, due to the additional means to finance it. The composition of wealth over the life cycle changes as well: while the average stock of durables is essentially the same, the average holdings of financial assets shifts towards older households in the alternative economy. This shift is a combination of three forces: first, allowing for unsecured borrowing mechanically decreases the net worth of constrained, young households; second, for given preferences and income risk, a looser borrowing limit induces lower precautionary savings; and third, the presence of annuities with similarly impatient households induces higher asset holdings at old ages.

We find that when unsecured debt is available, the rebalancing effects are smaller and are present over a shorter period of life. Figure 11 shows that, for both types of shocks, rebalancing is smaller for young households and it remains different from zero 5 years less 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 of income shocks to consumption. Table 4 shows that also the level of transmission coefficients for transitory shocks is lower in the economy with unsecured debt at all ages, while the transmission of permanent shocks is slightly higher. Hence, unsecured debt increases the overall ability to smooth shocks, but more so for transitory shocks. This is in contrast with the previous exercise, where an increase in the availability of collateralized debt led to lower rebalancing for both shocks, but it did not change much the level of transmission for nondurables. The fourth column in Table 4 shows how rebalancing translates into the difference between transmission to nondurable consumption and insurance. The overall bias decreases from -2.8 to -1.5 in the case of transitory shocks, and from 1.8 to 1.4 in the case of permanent shocks.

Figure 11: Difference in Transmission Coefficients (Unsecured Debt)

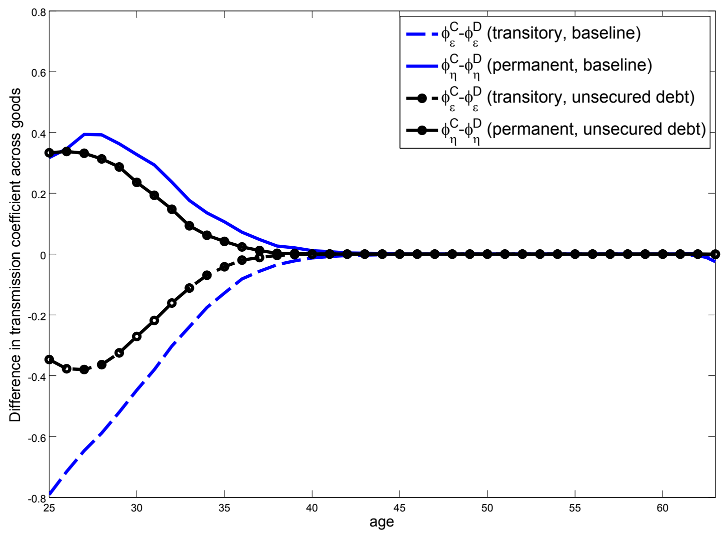

6 Conclusions

In this paper, we have analyzed the responses to income shocks of households that care for both durable and nondurable 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 assess 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 nondurable 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 shift of consumption towards one of the goods, depending on the persistence of the shock. When the shock is transitory, nondurable consumption reacts less than durable consumption, whereas the opposite is true when the shock is permanent.

Second, we have shown that insurance, defined as the ability to smooth a comprehensive measure of consumption across states, is a function of the transmission of income shocks to nondurable consumption and the extent of rebalancing. Therefore, the response to shocks of nondurable 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 larger for transitory than for permanent shocks. In our baseline calibration, the difference between the transmission of shocks to nondurable consumption and insurance was of the order of 3 percentage points for transitory shocks and 2 percentage points for permanent shocks. Among the young, this difference was 8 and 5 percentage points respectively.

We conducted a series of robustness checks that confirmed the role of rebalancing in consumption insurance across different model specifications. These exercises delivered some additional results. In particular, 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. We also found that the size of the bias caused by measuring insurance as the transmission to nondurable consumption alone depends on the required down payment on, and the durability of the other good, although the transmission itself does not. Finally, we found that the availability of uncollateralized loans matters both for the level of transmission and the age distribution of constrained households, and hence for the incidence of rebalancing over the life cycle. These two exercises combined imply that not only the level of credit available, but also its type (either collateralized or unsecured) is important to understand the size of the responses of nondurable consumption to income shocks and its accuracy as a measure of insurance.

As a final note, it would be interesting to test empirically whether unexpected income changes drive responses in the ratio of nondurable goods and the *stock* of durable goods that are different for constrained and unconstrained households. However, testing the rebalancing effect is notoriously difficult because of lack of good data on durable *stocks*. Recent empirical work finds evidence of asymmetric responses of nondurable consumption and *expenditure* of durable goods. For instance, Browning and Crossley (2009) show that among 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 (2013) 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 nondurable goods. Aaronson, Agarwal, and French (2012) find that households affected by a minimum wage hike increase expenditure on durables much more than in nondurables, 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. Hence, evidence of durable expenditure reacting more to transitory income shocks than nondurable expenditure does not need to reflect a rebalancing of durable and nondurable goods.

References

Aaronson, D., S. Agarwal, and E. French (2012): “The Spending and Debt Response to Minimum Wage Hikes,” *American Economic Review*, 102(7), 3111–3139.

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(6), 1169–1192.

Campbell, J., and J. Cocco (2003): “Household Risk Management and Optimal Mortgage Choice,” *Quarterly Journal of Economics*, 118(4), 1449–1494.

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.

Cocco, J. (2004): “Portfolio Choice in the Presence of Housing,” *Review of Financial Studies*, 18(2), 535–567.

Díaz, A., and M. J. Luengo-Prado (2010): “The Wealth Distribution with Durable Goods,” *International Economic Review*, 51(1), 143–170.

Fernández-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.

——— (2011): “Consumption And Saving Over The Life Cycle: How Important Are Consumer Durables?,” *Macroeconomic Dynamics*, 5(15), 725–770.

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.

HUGGETT, M. (1996): "Wealth Distribution in Life-Cycle Economies," *Journal of Monetary Economics*, 38(3), 469–494.

JOHNSON, D., R. MCCLELLAND, AND J. PARKER (2013): "Consumer Spending and the Economic Stimulus Payments of 2008," *American Economic Review*, 103(6), 2530–2553.

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

——— (2014): "A Model of the Consumption Response to Fiscal Stimulus Payments," forthcoming *Econometrica*.

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

MCBS PROJECT (2006): *Health & Health Care of the Medicare Population: Data from the 2002 Medicare Current Beneficiary Survey*. Westat Inc, Rockville, MD.

NAKAJIMA, M., AND I. TELYUKOVA (2012): "Home Equity in Retirement," mimeo.

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.

——— (2004): "Consumption and Risk Sharing over the Life Cycle," *Journal of Monetary Economics*, 51(3), 609–633.

YANG, F. (2009): "Consumption over the Life Cycle: How Different is Housing?," *Review of Economic Dynamics*, 12(3), 423–443.