


DOCUMENTOS DE TRABAJO

Consumption Insurance over the Life Cycle

Enzo A. Cerletti
Tomás Cortés

N° 1075 Marzo 2026

BANCO CENTRAL DE CHILE





La serie Documentos de Trabajo es una publicación del Banco Central de Chile que divulga los trabajos de investigación económica realizados por profesionales de esta institución o encargados por ella a terceros. El objetivo de la serie es aportar al debate temas relevantes y presentar nuevos enfoques en el análisis de los mismos. La difusión de los Documentos de Trabajo sólo intenta facilitar el intercambio de ideas y dar a conocer investigaciones, con carácter preliminar, para su discusión y comentarios.

La publicación de los Documentos de Trabajo no está sujeta a la aprobación previa de los miembros del Consejo del Banco Central de Chile. Tanto el contenido de los Documentos de Trabajo como también los análisis y conclusiones que de ellos se deriven, son de exclusiva responsabilidad de su o sus autores y no reflejan necesariamente la opinión del Banco Central de Chile o de sus Consejeros.

The Working Papers series of the Central Bank of Chile disseminates economic research conducted by Central Bank staff or third parties under the sponsorship of the Bank. The purpose of the series is to contribute to the discussion of relevant issues and develop new analytical or empirical approaches in their analyses. The only aim of the Working Papers is to disseminate preliminary research for its discussion and comments.

Publication of Working Papers is not subject to previous approval by the members of the Board of the Central Bank. The views and conclusions presented in the papers are exclusively those of the author(s) and do not necessarily reflect the position of the Central Bank of Chile or of the Board members.

Consumption Insurance over the Life Cycle*

Enzo A. Cerletti & Tomás Cortés
Central Bank of Chile

Resumen

En este trabajo estudiamos la capacidad de suavizado de consumo de los hogares a lo largo del ciclo vital. Para ello, aprovechamos al Encuesta Continua de Presupuestos Familiares española, que recolectó un panel con información detallada de consumo e ingreso con frecuencia trimestral entre 1985 y 1996. Documentamos la dinámica conjunta del consumo y el ingreso de los hogares españoles a lo largo del ciclo vital. Encontramos que la habilidad de los hogares para suavizar shocks de ingreso cambia drásticamente con la edad. En particular, la transmisión de shocks permanentes al consumo se reduce con la edad, desde una transmisión completa a los 25 años a menos del 40% a los 60. En cambio, la transmisión de shocks transitorios es baja a lo largo de todo el ciclo vital. Nuestros resultados son consistentes con modelos canónicos de ciclo vital con mercados incompletos, donde este patrón de comportamiento resulta de una combinación de una voluntad y capacidad de suavizado de shocks persistentes crecientes en la edad. Ilustramos este resultado calibrando un modelo cuantitativo para España. A su vez, validamos la generalidad de nuestro resultado empírico repitiendo la estimación con datos de Estados Unidos.

Abstract

This paper studies the amount of consumption insurance attained by households over the life cycle. We take advantage of the Spanish Household Budget Survey, which collected detailed panel data on consumption and income at the quarterly frequency between 1985 and 1996. We document the joint dynamics of consumption and income for Spanish households over the life cycle. We find that the ability of households to smooth income shocks changes sharply with age. In particular, the transmission of permanent shocks to consumption decreases with age, from full transmission at 25 to less than 40% at 60. In turn, the transmission of transitory shocks is low throughout the life cycle. Our results are broadly consistent with canonical life-cycle models of self-insurance, where the pattern emerges from the combination of increasing willingness and ability to smooth persistent shocks as households age. We illustrate this result by calibrating a quantitative model to Spain. We also validate the generality of our empirical results using data for the US.

* The views expressed are those of the authors and do not necessarily represent the views of the Central Bank of Chile or its board members. We thank Pedro Albarrán, Stéphane Bonhomme, Carlos Conesa, Eric French, Luisa Fuster, Michael Haliassos, María José Luengo-Prado, Claudio Michelacci and Josep Pijoan-Mas for helpful comments and discussions. All remaining errors are our own. An earlier version of this project circulated under the title "Life-cycle Patterns in Consumption Insurance: Evidence from the Spanish Household Budget Survey".
E-mail: ecerletti@bcentral.cl.

1 Introduction

How does the response of consumption to income shocks change over the life cycle? The most established theories of consumption with incomplete markets entail stark predictions for the transmission of income shocks into consumption at different points of the life cycle, yet direct empirical evidence on these predictions is scarce. In this paper, we attempt to fill this gap by providing empirical support to the predictions of a simple heterogeneous agents model that lies at the core of more sophisticated models. In such a model, the predicted age profile of the response of consumption to income changes is shaped by two fundamental forces: the *ability to smooth* stemming from self-insurance through savings, and the *willingness to smooth* based on the effective persistence of the shocks. In a finite horizon setup, no shock can be truly permanent: as the distance to the relevant horizon shortens, its effect on total lifetime resources converges to that of transitory shocks. While the existing literature has considered many refinements to the households' consumption and saving problem that would change the extent of self-insurance available at different ages, we argue that the willingness to smooth channel will still impose a marked decreasing profile for the response of consumption to permanent income shocks as households grow old, explaining why our empirical findings align with a rather stylized model.

One of the main building blocks of modern macroeconomics is household inter-temporal optimization under uncertainty. Following the seminal contributions of Modigliani and Brumberg (1954) and Friedman (1957), a long tradition in economics has studied the allocation of consumption and savings over the life cycle and its responses to changes in income. A crucial element for the theoretical analysis of consumption and income comovements is the set of securities that is available to consumers. In particular, in the extreme case of *complete markets*, a full set of Arrow-Debreu securities provides the ability to fully insure consumption against any sort of income fluctuations. On the contrary, if markets are *incomplete*, that is to say, if at least some securities cannot be traded, consumption will react to some extent to income innovations. A particularly relevant benchmark is the case in which consumers can only trade non-contingent assets, effectively being limited to self-insurance against income fluctuations. In this case, the consumption response to income changes depends critically on the persistence of income shocks, since transitory shocks are easily smoothed through borrowing and saving whereas very persistent shocks are not.

Most likely, the actual set of consumption insurance mechanisms available to households lies in between the cases of full insurance and self-insurance. In a very influential paper, Blundell et al. (2008) developed an empirical measure for the degree of consump-

tion insurance with respect to income shocks of different persistence. The extreme cases of complete markets and self-insurance can be characterized in terms of their predictions for the *transmission coefficients* of transitory and permanent income shocks into consumption. Blundell et al. (2008) estimate the transmission coefficients for the US economy and find evidence of full insurance with respect to transitory shocks (approximately consistent with both extreme cases) and partial insurance with respect to permanent shocks. The latter result implies that, while markets are indeed incomplete, there is some additional smoothing mechanism on top of self-insurance. The gap between these empirical findings and the predictions of a self-insurance model was explored by Kaplan and Violante (2010). They showed that a consumption model taking into account the life-cycle dimension of income uncertainty and consumption preferences predicts a level of transmission of permanent shocks which is lower than the canonical Permanent Income Hypothesis, but still higher than the estimates of Blundell et al. (2008). They conclude that the basic life-cycle, incomplete markets model is missing some additional mechanism for consumption smoothing.

The previous discussion revolves around the overall level of consumption insurance against idiosyncratic income shocks in the economy. However, life-cycle models of self-insurance also have strong predictions about how the transmission of income shocks into consumption changes with age. Specifically, in the basic model, the transmission of permanent income shocks to consumption declines steadily with age, while the transmission of transitory shocks remains low throughout the life cycle. Two forces shape this result. First, as pointed by Kaplan and Violante (2010), the gradual accumulation of wealth over the life cycle implies an increasing *ability* to smooth shocks. Second, as households approach retirement, permanent (labor) income shocks become effectively less persistent, insofar each individual shock becomes progressively less relevant for total expected lifetime income. Thus, households have an increasing *willingness* to smooth permanent shocks out. In this paper, we test this prediction in the data. Using the Spanish Household Budget Survey (*Encuesta Continua de Presupuestos Familiares, ECPF*), we are able to estimate the transmission of transitory and permanent income shocks to consumption at different stages of the life cycle. The age profile for consumption insurance that we obtain is broadly consistent with the predictions of the self-insurance model. In particular, there is a steep decline in the response of consumption to permanent shocks, but a flat profile for the consumption response to transitory shocks. Moreover, we find full transmission of permanent shocks into consumption for the youngest households, but less than 40% transmission for households close to retirement. The transmission for transitory shocks is indistinguishable from zero for all age groups. These empirical findings suggest

that self-insurance can explain an important part of the ability of households to smooth income shocks. Any alternative or additional channel for consumption smoothing that we might consider in order to account for the observed level of transmission of income shocks should be consistent with the marked age profile of consumption insurance as well.

While, to the best of our knowledge, no other paper has estimated a life-cycle profile for the transmission of income shocks into consumption, there is a literature exploring important deviations from the canonical model of self-insurance. Karahan and Ozkan (2013) consider an age-specific income process. Using US data, they find that persistent shocks are only moderately persistent at the beginning of the life cycle, gradually increasing to almost full persistence at age 45 and declining slightly afterwards. Its variance follows an opposite, U-shaped pattern. Bringing their estimates to a life-cycle model, they show that the transmission of permanent shocks into consumption is smaller for younger households, relative to an benchmark with an age-invariant income process, particularly when unsecured borrowing is allowed. In terms of our previous discussion, their results reflect a increased willingness to smooth shocks at young ages, due to the reduced persistence of shocks. Allowing for unsecured credit strengthens the result through a higher ability to smooth shocks for the young. In our model, we also allow for age-varying variances for the shocks, although we restrict our attention to an income process characterized by fully permanent and fully transitory shocks, which provides a good representation of observed income growth patterns over the life cycle, as shown by Hryshko (2012). Our empirical results on consumption insurance provide a benchmark on the net effect of both forces (willingness and ability to smooth) over the life cycle. In another influential paper, Kaplan and Violante (2014) present a model in which a proportion of middle-age households have a restricted ability to smooth shocks due to a lack of liquid assets, despite holding substantial wealth (the so-called wealthy hand-to-mouth). While omitted from our quantitative model, our empirical estimates should indirectly reflect the presence and relevance of these constrained agents in the Spanish data. More recently, Arellano et al. (2017) provide evidence of richer, nonlinear and history-dependent income dynamics, which lead to widespread heterogeneity in consumption responses to income fluctuations. In this paper, we restrict our attention to the canonical, linear income process and characterize the average transmission of income shocks into consumption for several age groups. Exploring the additional heterogeneity in consumption insurance within age groups documented by Arellano et al. (2017) is an interesting avenue for future research. Recent work by Ghosh and Theloudis (2025) has made progress in this direction.

The ECPF has been previously used to study the joint dynamics of income and consumption. Pijoan-Mas and Sánchez-Marcos (2010) take advantage of the longitudinal

dimension of the ECPF to study the evolution of income and consumption inequality in Spain during the 80s and the 90s. Albarrán et al. (2007) exploit the theoretical links between income and consumption growth rates to estimate earnings inequality in the ECPF. However, their focus is on comparing the relative importance of transitory and permanent shocks for employees and self-employed workers. Moreover, they impose full transmission of permanent shocks, and estimate the transmission coefficient of transitory shocks. The closest work to the present paper is Casado (2011), who replicates the study of Blundell et al. (2008) using the ECPF. In addition, he compares the results obtained using raw and imputed consumption data, finding a sizable downward bias in the level of insurance (an upward bias in the level of transmission) of permanent income shocks when using imputed data. Relative to his work, we focus on the life-cycle component of the transmission of income shocks to consumption. we also simulate a life-cycle model of consumption and savings calibrated to Spanish data, in order to compare its predictions to our empirical findings.

Finally, we explore the generality of our results by estimating a life-cycle profile for consumption insurance in the US. We take advantage of the information on both income and consumption collected in the Panel Study of Income Dynamics (PSID) survey since 1999 to overcome the need to impute consumption in older PSID data. These data have already been used by Blundell et al. (2016) to study the role of labor supply decisions as a consumption insurance channel. The drawbacks of the PSID data compared to the ECPF data is the smaller sampler size and its bi-annual frequency. The latter can be addressed through by adjusting our empirical specification to account for different time series properties compared to the quarterly Spanish data. In turn, the limited sample size makes it more challenging to construct an age profile for consumption insurance, which requires splitting the sample in smaller groups. To overcome this challenge, we employ a QMLE estimator, shown by Chatterjee et al. (2021) to yield tighter estimates than the minimum distance estimator in Blundell et al. (2008), and applied by us to the Spanish data, when conducting estimations in subgroups of the PSID.

The rest of the paper is organized in the following way. Section 2 presents the data and details the estimation methodology, while Section 3 discusses the estimation results. Then, a life-cycle model of consumption and savings is outlined in Section 4 in order to interpret the empirical results. Section 5 uses simplified version of the model to discuss analytically the key channels shaping consumption insurance over the life cycle. Finally, Section 6 concludes.

2 Data and empirical strategy

2.1 Data

We use the Spanish *Encuesta Continua de Presupuestos Familiares* (ECPF hereafter) to study empirically the transmission of income shocks to consumption. The ECPF is a quarterly rotating panel of households, each of which remains in the sample for at most 8 quarters. It collects detailed information on expenditures and income, as well as employment and demographic characteristics of the household. Appendix A provides a detailed definition of the income and consumption variables used for our analysis. The sample we use resembles the one in Pijoan-Mas and Sánchez-Marcos (2010), to whom we refer the reader for further details on the dataset. We performed some adjustments on their sample to match the sample selection criteria in Blundell et al. (2008). In particular, our sample spans from 1985 to 1996, and we focus in continuously married, working age individuals. We use quarterly monetary variables, deflated by the CPI to measure them in 1995 currency units, and transformed to euros using the fixed peseta-euro exchange rate set in December 31, 1998.

The advantage of using the ECPF is twofold. First, its length and sample size make it suitable for the estimation methodology followed. The identification strategy outlined in Meghir and Pistaferri (2004) requires at least three consecutive observations in income growth (hence four consecutive observations of income levels) to separately identify the variance of permanent and transitory shocks, as well as the transmission of each type of shock to consumption growth. More than 80% of the 148,679 household-quarter observations correspond to households observed for at least 4 consecutive quarters, and 47% of those correspond to households that stay in the sample for the maximum period of 8 quarters. The remaining observations are used to increase the precision in the estimates of consumption and income moments. On the other hand, the large cross-sectional dimension of the data allows us to conduct several exercises with different sub-samples. Of particular interest for the purpose of this paper is the fact that we can construct several age groups (six in the main exercise) and still have enough observations per cell to conduct a meaningful estimation.¹

Second, the ECPF contains detailed information on income and consumption at the quarterly frequency. This is in contrast with the lack of income information in the Consumer Expenditure Survey (CEX), where income is observed at the annual level, and the

¹We imposed a minimum of 50 observations per group for any sub-sample estimation, which cannot be attained in the imputed 1978-1992 PSID data for the same age 5-year age groups we use. By contrast, in our main exercise, the smallest group-quarter cell has 131 observations.

lack of information on consumption expenditures in the Panel Study of Income Dynamics (PSID), where only expenditures in food are present consistently throughout the sample, whereas more comprehensive consumption data is only available at biennial frequency since 1999. This means that the ECPF can be used directly for the purpose of estimating the transmission coefficients of income shocks to consumption, avoiding the imputation procedure followed by Blundell et al. (2008) to combine the information in the CEX and the PSID. Since the cross-sectional dimension of the PSID is much smaller than the one of the CEX, the imputed data set is not sufficiently large in order to perform the estimation at the age group level, as it is the goal of this paper. Moreover, Casado (2011) applies the imputation procedure in Blundell et al. (2008) to the ECPF and finds that the use of imputed data leads to upward-biased estimates of the transmission of permanent shocks to consumption. Hence, avoiding such imputations is an important advantage of the ECPF for the present empirical exercise, since it involves the estimation of several transmission coefficients.

2.2 Sample selection

We follow the sample selection criteria in Blundell et al. (2008) as closely as possible. Therefore, we focus on continuously married couples of working age (30 to 65 years old).² We drop income growth outliers, defined as those with changes in log labor income greater than 2 in absolute value.

Table 1: Sample descriptive statistics

Variable	Mean	s.d.
Age	47.37	10.25
Family size	4.07	1.37
N. of children	1.20	1.17
HS dropout or less	0.80	0.40
HS graduate or more	0.20	0.40
N. of labor income earners	1.34	0.82
Disposable income	3,813.20	2,367.99
Non-durable consumption	3,843.48	2,381.29
N. of observations	78,682	78,682

Note: Income and consumption figures are expressed in 1995 euros.

²Household heads below 30 represent 5% of the full sample. As in Blundell et al. (2008), we exclude them as an attempt to isolate our estimation from education and household formation decisions, which could be an important factor for self-selection into our sample. See Kaplan (2012) for a discussion on household formation decisions at young ages.

Our final sample comprises 78,682 household-quarter observations, obtained from 17,231 unique households. Regarding identification, 84.9% (62.5%) of the observations (households) correspond to households observed for at least four quarters, whereas 36.2% (21.1%) of the observations (households) come from individuals observed for the full 8-quarter spell. We lose one period when computing first-differences, leaving us with 75,924 household-quarter observations in terms of growth rates. Table 1 summarizes the main characteristics of our final sample. We also construct 6 groups define by age, at intervals of 5 years, between the ages of 30 and 60.³ Table 2 presents a comparison of means across these groups.

Table 2: Age groups: comparison of means

	30-34	35-39	40-44	44-49	50-54	55-59
Age	32.13	36.99	41.97	47.01	52.00	57.04
Family size	3.68	4.21	4.56	4.62	4.40	3.83
N. of children	1.59	2.07	2.06	1.37	0.75	0.38
HS dropout or less	0.65	0.68	0.76	0.82	0.85	0.90
HS graduate or more	0.35	0.32	0.24	0.18	0.15	0.10
N. of labor income earners	1.31	1.27	1.32	1.50	1.61	1.44
Disposable income	3,627.93	3,802.34	3,849.83	4,068.88	4,184.19	3,838.46
Non-durable consumption	3,343.16	3,786.69	3,980.31	4,301.82	4,317.90	3,857.17
N. of observations	10,152	11,706	11,868	10,995	10,605	10,660

Note: Income and consumption figures are expressed in 1995 euros.

The variables in Table 2 are distributed with an expected life-cycle pattern. Income grows and peaks at the early 50s, before early retirement becomes reduces average earnings. Consumption tracks the same hump-shaped pattern. Family size increases up to the late 40s, whereas the number of children peaks in the early 40s, and declines sharply thereafter. Educational attainment increases across cohorts, meaning that the share of households with high educational attainment declines with age. The overall level of education in the ECPF is consistent with the results in Fernández-Macías et al. (2013) for cohorts born between 1943 and 1989.⁴

³The 30 to 60 years old range is the most natural benchmark for our model, where we abstract from early retirement decisions, a more relevant phenomenon for workers in the 60 to 65 group. We keep all households between 30 and 65 for the full sample estimation, though.

⁴The oldest individuals in our final sample were born in 1926, while the youngest were born in 1966. High school enrollment and completion increased substantially in Spain during the second half of the 20th century, with the most drastic change involving those born between 1943 and 1963.

2.3 Estimation procedure

The empirical approach follows closely the one described in Blundell et al. (2008). We compute the unexpected growth of income and non-durable consumption as the first difference of the residuals of a regression of (log) income and consumption on a common set of observables. These regressors are meant to capture the predictable component of income and consumption growth, and include family size, number of earners, education, year dummies, seasonal dummies, a retirement status indicator, and several interactions. Then we estimate a variance-covariance matrix of the residuals for each quarter and stack those estimates in a vector m , which is fed to a minimum distance estimator of the form

$$\min_{\Lambda} (m - f(\Lambda))' A (m - f(\Lambda)) \quad (1)$$

where Λ is a vector of parameters, $f(\cdot)$ is a model-based mapping from Λ to the variances and covariances of income and consumption, and A is a weighting matrix. We set $A = V^{-1}$, where V is a diagonal matrix containing the variances of the estimates in m .

The vector function $f(\cdot)$ is obtained as an approximation of the solution to the model presented in the next section. We follow the non-durable consumption growth approximation in Blundell et al. (2007) to obtain an expression in terms of income shocks and transmission coefficients.

Consider the following income process, characterized by a deterministic component, a permanent component, and a purely transitory component:

$$\begin{aligned} \ln(Y_t) = \tilde{y}_t &= \mu_t + z_t + \varepsilon_t \\ z_t &= z_{t-1} + \eta_t \\ \varepsilon_t &\sim \mathcal{N}(0, \sigma_\varepsilon) \\ \eta_t &\sim \mathcal{N}(0, \sigma_\eta) \\ z_0 &\sim \mathcal{N}(0, \sigma_{z_0}) \end{aligned} \quad (2)$$

It follows that unpredictable income growth, defined as $\Delta y_t = \Delta \tilde{y}_t - \Delta \mu_t$, can be written as:

$$\Delta y_{it} = \eta_{it} + \varepsilon_{it} - \varepsilon_{it-1} \quad (3)$$

where η_{it} is an i.i.d. shock to the permanent component and ε_{it} is an i.i.d. transitory shock to income. The variance and covariances of income growth are given by

$$Var(\Delta y_{it}) = \sigma_{\eta,t}^2 + \sigma_{\varepsilon,t}^2 + \sigma_{\varepsilon,t-1}^2. \quad (4)$$

$$Cov(\Delta y_{it}, \Delta y_{it+1}) = -\sigma_{\varepsilon,t}^2 \quad (5)$$

where $\sigma_{x,t}^2$ denotes the variance of the variable x at period t . Notice that we allow for non-stationarity, i.e., the variances of the shocks are allowed to change over time.

A similar expression for consumption, denoted C_t , can be obtained for unconstrained households by combining an approximation of the Euler equation $u_c(C_t) = \beta(1+r_{t+1})E_t\{u_c(C_{t+1})\}$ with an approximation for the lifetime budget constraint, as described in Blundell et al. (2007), yielding

$$\Delta c_{it} = \phi_{\eta}\eta_{it} + \phi_{\varepsilon}\varepsilon_{it} + \xi_{it} \quad (6)$$

where $c_t = \ln(C_t)$, ϕ_{η} is the transmission coefficient of permanent income shocks to consumption, ϕ_{ε} is the transmission coefficient of transitory income shocks to consumption, and ξ_{it} captures unforeseeable heterogeneity in consumption growth, which can be interpreted as shocks to preferences. Under this approximation, the variance of consumption growth is given by:

$$Var(\Delta c_{it}) = \phi_{\eta}^2\sigma_{\eta,t}^2 + \phi_{\varepsilon}^2\sigma_{\varepsilon,t}^2 + \sigma_{\xi}^2 \quad (7)$$

with no serial correlation at any lead and lag.⁵

These (theoretical) variances and covariances are the $f(\cdot)$ functions used for estimation. Hence, the vector Λ of parameters is

$$\Lambda = \begin{bmatrix} \phi_{\eta} \\ \phi_{\varepsilon} \\ \sigma_{\xi}^2 \\ \sigma_{\eta}^2 \\ \sigma_{\varepsilon}^2 \end{bmatrix} \quad (8)$$

with σ_{η}^2 and σ_{ε}^2 are vectors containing the T values of $\sigma_{\eta,t}^2$ and $\sigma_{\varepsilon,t}^2$, respectively.⁶

⁵The presence of measurement error in the level of (log) consumption would cause the first autocovariance of consumption growth to be negative. The estimation procedure accounts for this possibility, estimating the implied variance of measurement error as well.

⁶In practice, it is not possible to identify T variances for each shock, since the first and the last observed periods lack the information (either past or future income growth) needed for identification. Moreover, in the last identified period, it is not possible to separate permanent and transitory shocks, requiring some additional restrictions. Hence, we estimate $T - 2$ variances of the transitory shock and $T - 4$ variances of the permanent shock.

3 Results

We conduct the estimation on the whole sample and also on subsamples defined by age. First, we describe the main features of income and consumption volatility in Spain. Then, we discuss the empirical transmission of income shocks to consumption, as captured by the transmission coefficients ϕ_η and ϕ_ε , with a special focus on how they vary with age.

3.1 Income and consumption in the ECPF


We begin the analysis by examining the volatility of income and consumption growth rates in the ECPF. Figure 1 illustrates the evolution of the volatility of the different variables of interest between 1985 and 1996. Panel (a) of Figure 1 shows that the variance of income growth, Δy_{it} , has declined mildly in the second half of the eighties, remaining stable during the nineties. The variation between quarters, however, is large relative to this trend. Consumption growth rates are much more volatile than its income counterparts, as shown in panel (b) of Figure 1, but also more stable over time. The variance of consumption growth does not exhibit the same decline as the variance of income. This is consistent with the findings of Pijoan-Mas and Sánchez-Marcos (2010), who study the evolution of inequality in Spain during the eighties and the nineties. They document a fall in income inequality during the time period covered by the ECPF, accompanied by a much smaller fall in consumption inequality.⁷ Panel (c) reveals that the co-movement of income and consumption has fallen steadily during the sample period covered by the ECPF. To some extent, this is in line with the relative decoupling of income and consumption inequality during the same period.

Panels (a) to (c) of Figure 1 contain the main components of the vector of moments m in (1).⁸ In turn, panel (d) shows the decomposition of the variance of income growth into the variance of permanent and transitory shocks obtained as an outcome of the minimum distance estimation. The most notable feature of this decomposition is the marked decline in the variance of permanent shocks between 1985 and 1992. Despite a moderate increase at the end of the sample, the relative importance of permanent income shocks is much lower in the nineties than it was in the eighties. Hence, the full picture of the evolution of income and consumption volatility during the period studied is characterized by declining income volatility, declining importance of permanent income shocks, stable consumption

⁷Pijoan-Mas and Sánchez-Marcos (2010) focus on the variance of the level of residual income to study the evolution inequality in Spain. In contrast, this paper estimates the variance of the first difference of residual income, as the focus here is the link between innovations to income and consumption growth rates.

⁸The remaining elements are covariances at different lags.

Figure 1: Income and consumption volatility over time



volatility, and decreasing correlation between income and consumption growth rates.

3.2 Transmission coefficients

Table 3 shows the results for overall consumption insurance under different restrictions on the variance of permanent shocks. The second column imposes no restriction, allowing the variance to be different each quarter of each year. The transmission coefficient for permanent shocks, ϕ_η , is 0.44 and significantly different from zero. This point estimate is lower than the one estimated by Blundell et al. (2008), but very similar to the one obtained by Casado (2011), suggesting that Spanish households are less sensitive to permanent income shocks than US households.⁹ On the other hand, the estimated transmission coefficient for transitory shocks, ϕ_ε , is of negligible size and statistically not different from zero. Therefore, full insurance with respect to transitory shocks cannot be rejected. This result is also in line with Casado (2011) and Blundell et al. (2008), who find no response to transitory shocks in Spain nor the US, respectively. The third column shows the results when imposing a polynomial structure to the variance of permanent shocks over time, trying to capture any low-frequency movements that may have taken place in Spain during the 80s and 90s. Finally, the fourth column shows the results when imposing seasonal structure, so that the variance of permanent shocks has a quarterly component, common across all years in the sample, and a year-specific component that varies freely. The results are very similar in all three cases. In particular, the seasonal model is almost identical to the unrestricted model, suggesting it sufficiently flexible to accommodate the patterns observed in the data, while having an advantage in terms of efficiency.

The main results from the estimation of (1) are summarized in Table 4, which reports the transmission coefficients for six age groups, from 30-35 to 56-60. The transmission of permanent shocks into consumption is positive and significant for all age groups. There are, however, large differences in the point estimates for different age groups. In particular, older households tend to exhibit much smaller responses to permanent shocks than households below 40. This pattern can be better appreciated in Figure 2. Panel (a) in Figure 2 shows the estimated transmission coefficients of permanent shocks to non-durable consumption for the six age groups. There appears to be a marked decline over the life

⁹The transmission coefficients estimated in the US are based on the comovements of annual consumption and annual income, while in the case of Spain quarterly data is used. Under the assumptions that justify (6), the frequency at which transmission is computed plays no role, so both estimates are comparable, leading to the conclusion that Spanish households are better insured. However, if the reduced-form model is misspecified, the different frequency of the datasets could be behind the lower figure for Spain. For instance, this would be the case if households need some months to incorporate changes in income into their consumption decisions.

Table 3: Estimation results: whole sample

Parameter	Unrestricted	Polynomial	Seasonal
ϕ_η	0.4675	0.4163	0.4612
(s.e.)	(0.0370)	(0.0366)	(0.0374)
ϕ_ε	-0.0493	-0.0368	-0.0459
(s.e.)	(0.0155)	(0.0157)	(0.0155)
N. of observations	75,924	75,924	75,924

Note: standard errors in parentheses.

cycle. At younger ages, the coefficient is indistinguishable from 1, meaning a full transmission of permanent shocks to income, as in the basic Permanent Income Hypothesis model. However, after age 50, less than 40% of a permanent shock passes on to consumption. At least two forces mitigate the impact of permanent income shocks in consumption with respect to the PIH benchmark. First, the durability of a permanent shock to labor income decreases mechanically as retirement age approaches. In the limit, a shock in the last quarter of an individual's working life lasts exactly one quarter regardless of its nature. Second, households accumulate assets over their life cycles, attaining a higher degree of self-insurance as their stock of assets increases relative to the size of the shocks faced. This is the basic result of a buffer-stock model of consumption and savings. Additional channels could be at work at higher ages in order to reduce the transmission of permanent income shocks to consumption. However, the high degree of transmission among the young points to a limited availability of alternative channels to smooth non-durable consumption.

Table 4: Estimation results: life cycle


Parameter	30-34	35-39	40-44	44-49	50-54	55-59
ϕ_η	1.0456	0.9298	0.5323	0.6148	0.3889	0.3836
(s.e.)	(0.1849)	(0.2093)	(0.0890)	(0.0884)	(0.0603)	(0.0686)
ϕ_ε	0.0356	-0.0101	-0.0189	-0.0599	0.0614	-0.0198
(s.e.)	(0.0544)	(0.0343)	(0.0397)	(0.0362)	(0.0380)	(0.0399)
N. of observations	7,566	9,234	9,385	8,610	8,304	8,326

Note: standard errors in parentheses.

In contrast, the transmission of transitory shocks to consumption is fairly small at all ages. Panel (b) in Figure 2 shows the age profile of the transmission coefficient of transitory shocks to non-durable consumption. Compared to the transmission of permanent shocks, the age profile of the transmission of transitory shocks looks flat around zero. In

fact, it is not possible to reject the null hypothesis of no transmission at all ages. This is in line with the basic PIH predictions, in the sense that the effect of transitory income shocks in consumption should be negligible as long as liquidity constraints are not too tight.

Figure 2: Age profiles of transmission coefficients




(a) Transmission coefficient with respect to permanent shocks (b) Transmission coefficient with respect to transitory shocks

Note: dashed lines represent 2 standard deviation bands

A similar life-cycle pattern for consumption insurance arises in the different specifications considered. Figure 3 shows the estimation results for broader age groups. Grouping households in 10 years intervals allows for more precise estimates, the downside being that each group is less homogeneous in terms of age composition. Again, we can see in panel (b) of Figure 3 that transitory shocks are fully insured at any stage of the life cycle. Panel (a), however, shows a clear decreasing pattern for the transmission of permanent income shocks to consumption. Compared to Figure 2, the slope is much flatter since, within each group, we are averaging across households of quite different ages.

The specification estimated in the first column of Table 4 is the most flexible in terms of the evolution of the variance of the shocks over time. As an alternative exercise, we imposed a polynomial structure for the variance of the permanent shock, aiming to capture its main trend in the 80s and 90s while reducing the number of parameters to be estimated. The estimation results under this different specification are in the second column of Table 4. It can be seen that the point estimates for the whole population are mostly insensitive to such changes in specification. However, the age profile of the transmission of permanent shocks to consumption is shifted somewhat when the evolution of the variance over time is restricted in this fashion. Figure 4 depicts the life-cycle profile of shock transmission

Figure 3: Age profiles of transmission coefficients: larger age groups



(a) Transmission coefficient with respect to permanent shocks

(b) Transmission coefficient with respect to transitory shocks


Note: dashed lines represent 2 standard deviation bands

in the restricted specification. We can see that the level of transmission of transitory shocks barely changes, but there is a large decrease in the transmission coefficients of permanent shocks, especially at young ages, where the estimation is more imprecise. It is not surprising that restricting the variance of permanent shocks has non-trivial effects on the estimated transmission coefficient, since both elements are tightly linked by (7) in the estimation procedure. We also tried imposing a seasonal structure on the variance of the permanent shock, allowing an annual component to vary freely throughout the sample period. The third column of Table 4 shows that the overall transmission of shocks in this case is very similar to our baseline, unrestricted estimates. Moreover, the age profile of the transmission of permanent shocks estimated with seasonality in the variance of shocks lies in between the unrestricted and the polynomial case (blue line in 4). Despite the difference in the exact level of transmission between specifications, the overall declining life-cycle pattern is present in all of them, appearing to be a robust feature of the data.

3.3 Life-cycle profiles of consumption insurance in the US

We provide further evidence on the generality of our results by replicating our findings in US data. Specifically, we use the Panel Study of Income Dynamics (PSID), which started collecting a comprehensive measure of consumption expenditures in 1999, at a biannual frequency. Using the post-1999 waves of the PSID overcomes the need to impute consumption expenditures, as in Blundell et al. (2008). However, the biannual frequency

Figure 4: Age profiles of transmission coefficients: polynomial for the variance of permanent shock



(a) Transmission coefficient with respect to permanent shocks (b) Transmission coefficient with respect to transitory shocks

Note: dashed lines represent 2 standard deviation bands.

implies a slow growth of its longitudinal dimension. At the time of writing this paper, there are twelve waves available with detailed data on both consumption and income (1999-2021). This is sufficient to identify the key parameters of our empirical specification. We construct our income and consumption variables, and our estimation sample, mimicking the criteria applied by Blundell et al. (2008) to the 1978-1992 sample. Appendix B provides further details on how we treat the PSID data.

While the current PSID has better information on consumption than the pre-1999 waves, the sample size is still smaller than that of the ECPF, making it difficult to achieve precise results when partitioning the sample. Therefore, to accommodate the lower frequency and smaller size of the US data, we make some adjustments to the empirical setup. First, we augment the income process with an MA(1) structure for the transitory shock, as in Blundell et al. (2008). We found it to be relevant for the biannual income data as well, but not in the quarterly Spanish data. Second, we restrict the variance of both shocks to be constant over time, reducing the number of parameters to be estimated. Our emphasis is on separately identifying the transmission coefficients for age groups as narrowly defined as the data allows us, which already puts pressure on our sample size. Since we allow for different variances for each age group, we consider this restriction to be an acceptable trade-off. Given the biannual structure of the data, the requirement of at least 4 consecutive observations to identify the parameters of interest, and the need for an additional consecutive observation to also identify the MA(1) parameter, the smallest

age groups we can define span 9 years. We also consider 10-year groups in the 30 to 60 years old range, the focus of our estimation in the Spanish data.

Even with these adjustments, obtaining precise estimates for sub-samples is challenging, as illustrated by the top panel of table 5. In fact, for some groups it's not possible to reject that the transmission of permanent and transitory shocks to consumption is the same. Hence, to add robustness to the results, we re-estimate the US age profiles using the quasi-maximum likelihood (QMLE) estimator applied by Chatterjee et al. (2021) to the imputed 1978-1992 PSID data from Blundell et al. (2008). The bottom panel of table 5 shows that these estimates are more precise, especially for the transmission coefficients of transitory shocks.

Table 5: PSID results

DWMD	26-34	35-43	44-52	53-61	30-39	40-49	50-59	Total
ϕ_η	0.7221	0.3596	0.1933	0.2886	0.3076	0.2216	0.2324	0.3367
(s.e.)	(0.3179)	(0.1445)	(0.0917)	(0.0825)	(0.0771)	(0.0841)	(0.0930)	(0.0385)
ϕ_ε	-0.0446	-0.0010	0.0177	-0.1660	0.0006	0.0383	-0.0107	-0.0431
(s.e.)	(0.0786)	(0.0580)	(0.0743)	(0.1169)	(0.0692)	(0.0688)	(0.0787)	(0.0275)
QMLE								
ϕ_η	0.8146	0.5481	0.2357	0.2515	0.3817	0.2890	0.2539	0.3951
(s.e.)	(0.2622)	(0.2172)	(0.0905)	(0.0632)	(0.0777)	(0.0757)	(0.0780)	(0.0351)
ϕ_ε	-0.0083	-0.0103	0.0053	-0.0104	-0.0022	0.0063	-0.0066	-0.0200
(s.e.)	(0.0079)	(0.0088)	(0.0084)	(0.0087)	(0.0080)	(0.0089)	(0.0110)	(0.0030)
N. of observations	4,044	4,909	3,758	3,103	5,731	4,796	3,863	21,965


Note: DWMD results correspond to the diagonally-weighted minimum distance estimator described in section 2.3 applied to the 1999-2021 PSID data. QMLE results correspond to the quasi-maximum likelihood estimator in Chatterjee et al. (2021).

Figure 5 illustrates our findings. The left panel shows a null transmission of transitory shocks at all ages, and a downward slopping profile, though very imprecise, for the transmission of permanent shocks, using the same minimum distance estimator as in the ECPF. The QMLE estimator, shown in the right panel, yields a cleaner profile for the transmission of both types of shocks. The decline in the transmission of permanent shocks over the life cycle is sharper, but substantial uncertainty around the point estimates remains.

Taken together, our results for Spain and the US are indicative of a marked decline in the transmission of permanent income shocks to consumption over the life cycle, although a precise quantitative estimation of its exact shape seems too demanding for the data

available.

Figure 5: Age profiles of transmission coefficients in the US



(a) 9-year groups, DWMD

(b) 9-year groups, QMLE

Note: dashed lines represent 2 standard deviation bands. DWMD results correspond to the diagonally-weighted minimum distance estimator described in section 2.3 applied to the 1999-2021 PSID data. QMLE results correspond to the quasi-maximum likelihood estimator in Chatterjee et al. (2021).

4 Model

4.1 A life-cycle model of consumption and savings

This section outlines a simple model of consumption and savings with transitory and permanent income shocks. This framework should allow us to understand the determinants of consumption responses to changes in income over the life cycle.

Consider a household that lives for a maximum of T periods. During the first T_R periods, it obtains a stochastic stream of labor income subject to permanent and transitory shocks, as described in equations (2)-(3), living as retirees thereafter. Its life-time utility, U , depends on the consumption sequence $\{C_t\}_{t=1}^T$ as follows

$$U = \sum_{t=1}^T \beta^t u(C_t), \quad (9)$$

where $u(\cdot)$ is a period utility function and β is a discount factor capturing time preferences. The household faces a period budget constraint of the form

$$A_{t+1} + C_t = (1 + r)A_t + Y_t, \quad (10)$$

where A_t is a stock of assets and Y_t is a stochastic income flow. Assets are restricted to be above a threshold $\underline{A}$:

$$A_{t+1} \geq \underline{A}. \quad (11)$$

Households maximize (9) subject to a sequence of T budget constraints of the form in (10) and T borrowing constraints of the form in (11). At each period t , optimality of the consumption decision requires

$$u_c(C_t) = \lambda_t, \quad (12)$$

where λ_t is the multiplier associated to the t^{th} period budget constraint. In turn, optimal savings satisfy

$$\lambda_t = E_t [\beta(1+r)\lambda_{t+1}] - \gamma_t, \quad (13)$$

where γ_t is the multiplier associated to the $t - th$ period borrowing constraint. Combining the two delivers the following intertemporal condition:

$$u_c(C_t) = E_t [\beta(1+r)u_c(C_{t+1})] - \gamma_t. \quad (14)$$

Whenever borrowing constraints are not binding, $\gamma_t = 0$ and a standard Euler equation arises, $u_c(C_t) = E_t [\beta(1+r)u_c(C_{t+1})]$. However, binding borrowing limits introduce a wedge in intertemporal allocations, measured by the multiplier γ_t . Hence, income shocks will transmit into consumption through two channels. Intuitively, income shocks change the level of income, Y_t , and also expectations about future income, depending on the persistence of the shock. Households will adjust their consumption level in response to news about their total lifetime resources. This is a standard permanent income channel. The second channel comes from the *timing*, rather than the level of income. By changing the sequence of household income, $\{Y_j\}_{j=t}^T$, an innovation to income can also change the intertemporal wedges, $\{\gamma_j\}_{j=t}^T$. While the permanent income channel always plays a role for all households, the relevance of the second channel will depend on whether a household is at or close to the liquidity constraint and the persistence of the income shock. Therefore, alternative assumptions about the borrowing capacity of households and the stochastic process for Y_t will have non-trivial effects on the links between income and consumption. The baseline configuration of the model features the income process described in (2), so the dynamics of the model are driven by the stochastic process estimated in Section 2. However, by imposing an exogenous borrowing limit $\underline{A}$, the model described in this section is not accurately approximated by the reduced-form model in section 2. As highlighted by Kaplan and Violante (2010), the reduced-form model is misspecified in the presence of binding borrowing limits.

4.2 Calibration

The model is calibrated following as closely as possible the structure of the ECPF. The baseline parameterization is summarized in Table 6. The time unit is a quarter, and households live for up to 280 quarters. We assume households are born as workers at the age of 25, making decisions from age 26 onward, and retire at 60, which corresponds to $t = 140$. They live as retirees for up to 140 additional quarters, with a decreasing survival probability. During retirement, households receive a constant pension benefit, which is a concave function of the average of their working life earnings.

Table 6: Calibrated parameters

<i>Manually set</i>	Value			
σ	2			
r	0.74%			
A	0			
<i>Estimated in the ECPF</i>	Value			
σ_ε^2 (26-35)	0.036			
σ_ε^2 (36-40)	0.042			
σ_ε^2 (41-45)	0.073			
σ_ε^2 (46-50)	0.068			
σ_ε^2 (51-55)	0.042			
σ_ε^2 (56-60)	0.044			
σ_η^2 (25-35)	0.020			
σ_η^2 (36-40)	0.015			
σ_η^2 (41-45)	0.021			
σ_η^2 (46-50)	0.024			
σ_η^2 (51-55)	0.032			
σ_η^2 (56-60)	0.025			
$\sigma_{z_0}^2$	0.148			
<i>Calibrated in equilibrium</i>	Value	Target	Model	Data
β	1.006	average W/Y (EFF)	21.61	21.61
<i>Note: W is defined as total net worth both in the data and in the model.</i>				

The parameters of the income process come from our estimates in section 2. Since the model has no aggregate dynamics, a single variance for each type of shock and age group is used. This is only problematic for the variance of permanent shocks, which follows a U shape between 1985 and 1996. As a first approximation, we feed the model the time-series average variance of each type of shock and age group, and then assess the sensitivity of the results to different levels of income uncertainty. The second panel in table 6 shows the average variance across age groups for each type of shock. The deterministic, age-

specific component of income, μ_t , is obtained from the regression of income on observables that delivered our measure of residual income growth. In particular, we re-scale the coefficients of the age dummies in that regression so as to replicate the average income in the sample. Finally, we generate an initial log-normal distribution for the permanent shock with variance σ_{z_0} . The variance of this initial shock is set so the model replicates the degree of income inequality at age 25 given the estimates for σ_ε^2 and σ_η^2 .

The interest rate is set to 0.74% per quarter, implying an annual interest rate of 3%. Then, the discount factor β is chosen to generate the average wealth to average income ratio in the Spanish Household Finance Survey (*Encuesta Financiera de las Familias, EFF*).¹⁰

The period utility function is assumed to be a standard CRRA function with coefficient of risk aversion σ :

$$u(C_t) = \frac{C_t^{1-\sigma} - 1}{1-\sigma}. \quad (15)$$

In the baseline calibration, σ is set to 2, as widely used in this type of models.


4.3 Simulation results

The model presented in section 4.1 is solved numerically. We use the resulting policy functions and a set of initial conditions to create a sample of simulated households. In the baseline calibration, the initial condition for the permanent shock, z_0 , is drawn from its calibrated distribution, while the initial condition for assets, A_1 , is set to zero for all households. Then, I use the simulated sample to compute the transmission coefficients with respect to each type of shock for the same age groups considered in Section 3.2. The simulated transmission coefficients are compared to the estimated age profiles. Thus, we can evaluate the extent to which the simple model can replicate the life-cycle profile for consumption insurance measured in Spanish data.


Figure 6 summarizes the results. Panels (a) and (c) show the transmission coefficients of permanent income shocks to consumption, for 5-year age groups and 10-year age groups respectively. Panels (b) and (d) depict the transmission on transitory shocks. It can readily be seen that the model (solid green line) easily generates the low response of consumption to transitory income shocks observed in the data, even for young households. On the other hand, the transmission coefficients of permanent shocks display a steep de-

¹⁰The EFF is a survey conducted by the Bank of Spain every 3 years since 2002. It collects detailed information on household portfolio allocations, as well as annual income and consumption expenditures. Hence, the wealth to annual income ratio obtained from the EFF is properly scaled to compare to the wealth to quarterly income ratio generated by the model. We use the first wave, which is closest in time to our consumption and income data.


Figure 6: Age profiles of transmission coefficients




(a) Transmission coefficient with respect to permanent shocks



(b) Transmission coefficient with respect to transitory shocks



(c) Transmission coefficient with respect to permanent shocks



(d) Transmission coefficient with respect to transitory shocks

cline over the life cycle, roughly consistent with its empirical counterparts. In fact, the model-generated age profile of the transmission coefficients of permanent shocks lies in between the empirical profiles obtained under different specifications. The most remarkable difference between the model and the data is perhaps the moment in life where the decline is steeper. While in the data the transmission coefficient seems to decline quickly between the ages of 30 and 45, the model delivers a consistent decline over the entire working life, accelerating somewhat after age 46.

In sum, the model does a good job in replicating the sensitivity of consumption to both permanent and transitory income shocks at different stages of the life-cycle. Nevertheless, it seems to miss the exact timing of the transition from high transmission to low transmission of permanent shocks during the life cycle. Although the precision of the empirical transmission coefficients does not allow us to reject the transmission level generated by the model at any age, it is worth studying extensions of the model that may change the particular shape of the age profile of transmission coefficients, such as the age-varying persistence of shocks proposed by Karahan and Ozkan (2013).

5 Discussion: the willingness to smooth over the life cycle

In the previous section, we established that a relatively simple model can account for the observed pattern of transmission of permanent shocks over the life cycle. What can account for the success of the relatively stripped down model? We argue that the empirical results point to a rather high level of self-insurance in Spain, so that changes in transmission over the life cycle are mostly driven by changes in the willingness to smooth shocks as households age. In this section, we show how this mechanism works by making additional simplifying assumptions on the model presented in section 4.

First, let us assume no uncertainty regarding lifespan or income, and no binding borrowing limit. Households live T periods, and we replace the income process with $Y_t + 1 = \rho Y$ for all $t < T_R < T$, and $Y_t = 0$ afterwards. This choice of income allows us to examine income "shocks" in the deterministic setup by looking at the link between consumption at time t to different values of Y_t taking into account the information it contains about future income. The retirement age T_R marks the point in the life cycle after which income no longer depends on Y_t . For simplicity, we also assume $\beta(1+r) = 1$. From (14) and (15), these assumptions imply households desire a flat consumption profile, $C_{t+1} = C_t \forall t$. To illustrate our point, we consider two versions of the simplified model: first, we consider an infinite horizon case, with $T_R = T \rightarrow \infty$ and $0 < \rho \leq 1$.¹¹ Second,

¹¹The infinite horizon version of the simplified model follows closely the discussion in section 2.10 of

we consider the finite horizon version with fully permanent shocks, so $T_R < T$ and $\rho = 1$.

Starting with the infinite horizon case, the intertemporal budget constraint would be

$$\sum_{j=0}^{\infty} (1+r)^{-j} C_t = \sum_{j=0}^{\infty} (1+r)^{-j} \rho^j Y_t + A_t \quad (16)$$

$$\Rightarrow \frac{1+r}{r} C_t = \frac{1+r}{1+r-\rho} Y_t + A_t \quad (17)$$

$$C_t = \frac{r}{1+r-\rho} Y_t + \frac{r}{1+r} A_t \quad (18)$$

Equation (18) shows the role of persistence in shaping the consumption response to additional resources. The derivative of current consumption with respect to current assets, $\frac{\partial C_t}{\partial A_t} = \frac{r}{1+r}$, is the intended transmission of transitory shocks when no liquidity constraint hinders self-insurance and we abstract from precautionary motives. For persistent shocks, the fraction of the shocks that households desire to translate into consumption (conversely, the fraction they do not want to smooth out) is $\frac{\partial C_t}{\partial Y_t} = \frac{r}{1+r-\rho}$, which is larger than for transitory shocks for $\rho > 0$ and converges to $\frac{r}{1+r}$ as $\rho \rightarrow 0$. For fully permanent shocks, the response is identical to 1, the standard Friedman (1957) result. Intermediate levels of transmission correspond to intermediate values of ρ .

Now we turn our attention to the finite horizon case with $\rho = 1$. The intertemporal budget constraint now is

$$\sum_{j=0}^{T-t} (1+r)^{-j} C_t = \sum_{j=0}^{T_R-t} (1+r)^{-j} Y_t + A_t \quad (19)$$

$$\Rightarrow C_t = \frac{\sum_{j=0}^{T_R-t} (1+r)^{-j}}{\sum_{j=0}^{T-t} (1+r)^{-j}} Y_t + \frac{A_t}{\sum_{j=0}^{T-t} (1+r)^{-j}} \quad (20)$$

We have the equivalent of (18) in (20), which shows the role of the distance to the labor income horizon, $T_R - t$, in shaping the response of consumption to changes in income. However, it is a convoluted expression to work with for the purpose of discussing the role of T_R . For the sake of clarity, we can additionally impose that $r = 0$ in the finite horizon case, yielding

$$C_t = \frac{T_R - t}{T - t} Y_t + \frac{A_t}{T - t} \quad (21)$$

Provided $T_R < T$, equation (21) shows that even permanent shocks will only partially

Cerletti and Pijoan-Mas (2012), excluding durable goods and with the additional assumption $\beta(1+r) = 1$.

translate into consumption, while transitory shocks are divided evenly across the remaining periods of life. Importantly, though, the passthrough from permanent shocks to consumption falls with t , as the households approaches T_R : while at sufficiently long horizons, $t \rightarrow -\infty$, it converges to the permanent income hypothesis value of one, at shorter horizons it falls towards the simple annuitization factor of transitory shocks. When $t = T_R - 1$, $\frac{\partial C_t}{\partial Y_t} = \frac{T_R - T_R + 1}{T - T_R + 1} = \frac{1}{T - T_R}$, which coincides exactly with $\frac{\partial C_t}{\partial A_t}$ at the same age. Hence, aging in a finite horizon model with fixed persistence of shocks has the same effect as reducing the persistence of shocks in an infinite horizon model when studying the response to shocks coming mostly from labor income.

This result is hardly novel, as it can be found in any basic cake-eating problem. However, we think it is important to highlight it in the context of consumption insurance over the life cycle, as it strongly constrains the shape of the age profile of transmission of income shocks that any model of self-insurance can attain. We stress that this declining willingness to smooth under incomplete markets comes entirely from household preferences and their desired, unconstrained response to income changes. Departures from our basic modeling of self-insurance that increase the response of older households to permanent shocks through a lower ability to insure would also imply stronger responses to transitory shocks, as in Kaplan and Violante (2014). Flattening the response of consumption at early stages of the life-cycle, however could not come from additional means of self-insurance, since young households are *unwilling* to smooth regardless of their ability to do so. Hence, any dampening of their response to permanent shock is more likely to come from mechanisms that partially restore market completeness.¹² We reiterate that, in the Spanish case, the empirical evidence is consistent with a basic setup without additional refinements. However, this result may hide the interaction of several offsetting mechanisms operating simultaneously, and may not be as strong in other economies. In particular, it could be worth exploring the interaction between distance to retirement and age-varying persistence of shocks as proposed by Karahan and Ozkan (2013).

6 Conclusions

One of the main predictions of all the life-cycle models of consumption and savings with incomplete markets is that the exposure of households to income risk should decrease with age, as they accumulate assets and become better self-insured. In this paper, we tested this prediction by estimating the transmission coefficient of permanent and transitory income

¹²Kaplan (2012) provides an example of a mechanism by which young workers may be insured beyond self-insurance through saving and borrowing.

shocks into consumption at different stages of the life cycle. Using a panel of Spanish households with quarterly information on both income and consumption, we found broad support for the hypothesis of decreasing transmission over the life cycle.

We also calibrated a standard life-cycle model of consumption and savings under income uncertainty to the Spanish economy, in order to check whether this widely used theoretical framework is consistent with the empirical profile of transmission coefficients. We found that the model, despite its simplicity, does a remarkable job in reproducing the main features of consumption insurance over the life cycle. We show that a simple mechanism at the core of the standard life-cycle model, governing the households' willingness to smooth shocks, is behind the decline of transmission coefficients to permanent shocks as households age, and that it is hard to overcome within the incomplete-markets framework by refinements to the households' ability to smooth them.

A caveat to the empirical analysis conducted in this paper is that its structural interpretation requires assuming the absence of binding liquidity constraints, while these are important driving forces of the age profile of transmission in the model. However, borrowing constraints are hard to identify in the data, and the same life-cycle profile of consumption insurance could be originated by alternative mechanisms. Therefore, a challenge for future research is to identify the incidence of borrowing constraints at different stages of the life cycle, in order to test a more precise prediction of the theoretical model. Nevertheless, the estimates presented here retain a less-structural interpretation in the presence of borrowing constraints or other deviations from the baseline model, as a linear projection of residual consumption growth onto residual income growth. Such projection can be performed in both actual and simulated data, as done in section 4.3, supporting the interpretation of our results.

We have explored the robustness of our findings to alternative estimation strategies and data sources, substantiating our findings as a general pattern in advanced economies. In particular, we extended our empirical work in two ways. First, we replicated our empirical exercise with US data, taking advantage of the richer information on consumption collected by the PSID since 1999. Second, since our empirical strategy involves partitioning the sample, and the PSID sample is smaller than the ECPF sample, improving the precision of our estimates is a primary concern. We implemented the estimator proposed by Chatterjee et al. (2021), who suggest an efficient alternative to the estimation procedure in Blundell et al. (2008). We found that US data supports the life-cycle profiles for consumption insurance estimated on Spanish data, although it is more challenging to obtain sharp estimates in it. Using QMLE helps somewhat, and further supports a flat profile for the transmission of transitory shocks.

References

- ALBARRÁN, P., R. CARRASCO, AND M. MARTÍNEZ-GRANADO (2007): “Inequality for Wage Earners and Self-Employed: Evidence from Panel Data,” Working Paper 07-24, Universidad Carlos III.
- ARELLANO, M., R. BLUNDELL, AND S. BONHOMME (2017): “Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework,” *Econometrica*, 85, 693–734.
- BLUNDELL, R., H. LOW, AND I. PRESTON (2007): “Income Risk and Consumption Inequality: A Simulation Study,” Institute for Fiscal Studies Working Paper 04/26.
- BLUNDELL, R., L. PISTAFERRI, AND I. PRESTON (2008): “Consumption Inequality and Partial Insurance,” *American Economic Review*, 98, 1887–1921.
- BLUNDELL, R., L. PISTAFERRI, AND I. SAPORTA-EKSTEN (2016): “Consumption Inequality and Family Labor Supply,” *American Economic Review*, 106, 387–435.
- CASADO, J. M. (2011): “From income to consumption: measuring households partial insurance,” *Empirical Economics*, 40, 471–495.
- CERLETTI, E. A. AND J. PIJOAN-MAS (2012): “Durable Goods, Borrowing Constraints and Consumption Insurance,” CEMFI Working Paper No. 1206.
- CHATTERJEE, A., J. MORLEY, AND A. SINGH (2021): “Estimating household consumption insurance,” *Journal of Applied Econometrics*, 36, 628–635.
- FEENBERG, D. AND E. COUTTS (1993): “An Introduction to the TAXSIM Model,” *Journal of Policy Analysis and Management*, 12, 189–194.
- FERNÁNDEZ-MACÍAS, E., J.-I. ANTÓN, F.-J. B. A, AND R. M. OZ DE BUSTILLO (2013): “Early School-leaving in Spain: evolution, intensity and determinants,” *European Journal of Education*, 48, 150–164.
- FRIEDMAN, M. (1957): *A Theory of the Consumption Function*, Princeton University Press.
- GHOSH, A. AND A. THELOUDIS (2025): “Consumption Partial Insurance in the Presence of Tail Income Risk,” Papers 2306.13208v5, arXiv.org.
- HRYSKO, D. (2012): “Labor income profiles are not heterogeneous: Evidence from income growth rates,” *Quantitative Economics*, 3, 177–209.
- KAPLAN, G. (2012): “Moving Back Home: Insurance Against Labor Market Risk,” *Journal of Political Economy*, 120, 446–512.
- KAPLAN, G. AND G. VIOLANTE (2010): “How Much Consumption Insurance Beyond Self-Insurance?” *American Economic Journal: Macroeconomics*, 2, 53–87.

- KAPLAN, G. AND G. L. VIOLANTE (2014): “A Model of the Consumption Response to Fiscal Stimulus Payments,” *Econometrica*, 82, 1199–1239.
- KARAHAN, F. AND S. OZKAN (2013): “On the persistence of income shocks over the life cycle: Evidence, theory, and implications,” *Review of Economic Dynamics*, 452–476.
- MEGHIR, C. AND L. PISTAFERRI (2004): “Income Variance Dynamics and Heterogeneity,” *Econometrica*, 72, 1–32.
- MODIGLIANI, F. AND R. H. BRUMBERG (1954): “Utility analysis and the consumption function: an interpretation of cross-section data,” in *Post-Keynesian Economics*, ed. by K. K. Kurihara, Rutgers University Press, 388–436.
- PIJOAN-MAS, J. AND V. SÁNCHEZ-MARCOS (2010): “Spain is Different: Falling Trends of Inequality,” *Review of Economic Dynamics*, 13, 154–178.

A Definitions in the ECPF data

Disposable income is defined as the total net non-financial income of the household. It includes wages, self-employment income, unemployment benefits, and pensions.¹³ Non-durable consumption is the sum of expenditures on several goods and services, comprising food, beverages and tobacco, clothing, utilities, cleaning products and services, medical expenditures, transportation fees and gasoline, entertainment, restaurants, and other non-durable goods and services.

B Definitions and sample selection in the PSID data

We use consumption and income data from the PSID for the period 1999-2021. Hence, we have 12 waves collected every 2-years, resulting in 11 biannual growth rates per household. We follow the definitions and sample restrictions in Blundell et al. (2008) as close as possible.

Our measure of income is total non-financial taxable income, which includes labor income, the labor part of business income, unemployment benefits, and other transfer income for all adult members of the household. We use TAXSIM v.35 to obtain after-tax income (see Feenberg and Coutts (1993)). Non-durable consumption is the sum of expenditures on food at home and outside, utilities, gasoline, parking, and public transportation fees. Income and consumption magnitudes are measured at the yearly level, and deflated by the CPI.

To remove the predictable component of income and consumption growth, we run separate regressions of income and consumption on year, region, race, age, retirement status, and education dummies, family size, number of children in the household, and interaction terms between year and all other variables. We use the residuals of these regressions as input for our estimation procedures.

¹³Households are asked to report after-tax earnings, so no adjustment is necessary to obtain net income.

Our sample consists of working-age continuously married couples, defined as households headed by an individual between 25 and 65 years old and no change in marital status. Moreover, we exclude households where the head is a student, a housewife, or permanently disabled. We drop observations with missing data for income or consumption. We also remove a few outliers, defined as households with less than 100 dollars of annual income or with extreme income growth rates.

C Additional estimation results

In this section, we report the main results of the paper together with the estimated variances of the income shocks over time. For brevity, we report the average variance of quarterly shocks for each year, while the full set of quarterly variances for the unrestricted case are depicted in panel (d) of Figure 1.

Table 7: Estimation results: whole sample

Parameter		Unrestricted	Polynomial	Seasonal
ϕ_η		0.4675	0.4163	0.4612
(s.e.)		(0.0370)	(0.0366)	(0.0374)
ϕ_ε		-0.0493	-0.0368	-0.0459
(s.e.)		(0.0155)	(0.0157)	(0.0155)
σ_η^2	1985	0.0364	0.0371	0.0359
	1986	0.0319	0.0325	0.0319
	1987	0.0279	0.0278	0.0273
	1988	0.0213	0.0237	0.0209
	1989	0.0182	0.0203	0.0185
	1990	0.0160	0.0175	0.0158
	1991	0.0181	0.0153	0.0180
	1992	0.0133	0.0138	0.0129
	1993	0.0154	0.0130	0.0149
	1994	0.0148	0.0128	0.0147
	1995	0.0114	0.0132	0.0109
	1996	0.0122	0.0142	0.0121
σ_ε^2	1985	0.0460	0.0455	0.0457
	1986	0.0364	0.0361	0.0365
	1987	0.0386	0.0390	0.0391
	1988	0.0443	0.0430	0.0440
	1989	0.0354	0.0348	0.0353
	1990	0.0350	0.0347	0.0329
	1991	0.0347	0.0356	0.0347
	1992	0.0380	0.0381	0.0383
	1993	0.0406	0.0413	0.0408
	1994	0.0383	0.0390	0.0385
	1995	0.0408	0.0393	0.0402
	1996	0.0372	0.0368	0.0372
N. of observations		75,924	75,924	75,924

Note: reported variances for each year are averages over the four quarterly estimates.

Table 8: Estimation results: life cycle

Parameter		30-34	35-39	40-44	44-49	50-54	55-59
ϕ_η		1.0456	0.9298	0.5323	0.6148	0.3889	0.3836
(s.e.)		(0.1849)	(0.2093)	(0.0890)	(0.0884)	(0.0603)	(0.0686)
ϕ_ε		0.0356	-0.0101	-0.0189	-0.0599	0.0614	-0.0198
(s.e.)		(0.0544)	(0.0343)	(0.0397)	(0.0362)	(0.0380)	(0.0399)
σ_η^2	1985	0.0371	0.0362	0.0340	0.0468	0.0653	0.0378
	1986	0.0293	0.0277	0.0316	0.0391	0.0546	0.0371
	1987	0.0220	0.0196	0.0289	0.0317	0.0440	0.0358
	1988	0.0164	0.0133	0.0263	0.0255	0.0353	0.0342
	1989	0.0126	0.0087	0.0238	0.0207	0.0284	0.0320
	1990	0.0105	0.0058	0.0214	0.0172	0.0233	0.0294
	1991	0.0102	0.0047	0.0191	0.0150	0.0200	0.0264
	1992	0.0116	0.0053	0.0169	0.0141	0.0185	0.0230
	1993	0.0147	0.0077	0.0148	0.0146	0.0189	0.0190
	1994	0.0196	0.0118	0.0128	0.0164	0.0211	0.0147
	1995	0.0263	0.0176	0.0108	0.0195	0.0251	0.0099
	1996	0.0347	0.0251	0.0090	0.0240	0.0310	0.0047
σ_ε^2	1985	0.0479	0.0489	0.0543	0.0392	0.0412	0.0649
	1986	0.0311	0.0365	0.0311	0.0415	0.0407	0.0336
	1987	0.0421	0.0437	0.0345	0.0428	0.0459	0.0373
	1988	0.0322	0.0341	0.0401	0.0373	0.0486	0.0549
	1989	0.0316	0.0328	0.0305	0.0331	0.0460	0.0383
	1990	0.0329	0.0383	0.0282	0.0373	0.0351	0.0430
	1991	0.0321	0.0456	0.0454	0.0325	0.0471	0.0280
	1992	0.0410	0.0301	0.0251	0.0435	0.0486	0.0402
	1993	0.0356	0.0494	0.0375	0.0335	0.0478	0.0420
	1994	0.0320	0.0424	0.0446	0.0390	0.0371	0.0483
	1995	0.0387	0.0536	0.4644	0.0500	0.0295	0.0473
	1996	0.0311	0.0440	0.0352	0.388	0.0387	0.0453
N. of observations		7,566	9,234	9,385	8,610	8,304	8,326

Note: reported variances for each year are averages over the four quarterly estimates.

Documentos de Trabajo Banco Central de Chile	Working Papers Central Bank of Chile
NÚMEROS ANTERIORES	PAST ISSUES
La serie de Documentos de Trabajo en versión PDF puede obtenerse gratis en la dirección electrónica: www.bcentral.cl/esp/estpub/estudios/dtbc .	Working Papers in PDF format can be downloaded free of charge from: www.bcentral.cl/eng/stdpub/studies/workingpaper .
Existe la posibilidad de solicitar una copia impresa con un costo de Ch\$500 si es dentro de Chile y US\$12 si es fuera de Chile. Las solicitudes se pueden hacer por fax: +56 2 26702231 o a través del correo electrónico: bcch@bcentral.cl .	Printed versions can be ordered individually for US\$12 per copy (for order inside Chile the charge is Ch\$500.) Orders can be placed by fax: +56 2 26702231 or by email: bcch@bcentral.cl .

DTBC – 1075

Consumption Insurance over the Life Cycle

Enzo Cerletti, Tomás Cortés

DTBC – 1074

Precios de viviendas en Chile: Herramientas para Evaluar Desalineamientos y sus Efectos sobre la Banca

Serio Díaz V., Mauricio Salas G., Francisco Vásquez L.

DTBC – 1073

The Life Experience of Central Bankers and Monetary Policy Decisions: A Cross-country Dataset

Carlos Madeira

DTBC – 1072

Coordinating in the Haircut. A Model of Sovereign Debt Restructuring in Secondary Markets

Adriana Cobas

DTBC – 1071

Liquidity Stress Tests for Fixed-Income Mutual Fund: an application for Chile

Tamara Gallardo, Fernando Martínez, Matías Muñoz, Félix Villatoro

DTBC – 1070

Climate Transition Risks in Chile's Banking Industry: A Loan-Level Stress Test

Felipe Córdova, Francisco Pinto, Mauricio Salas

DTBC – 1069

How accurately do consumers report their debts in household surveys?

Carlos Madeira

DTBC – 1068

Riesgo de Crédito Gestionado por Medio de un Modelo de Espacio-Estado Aplicado a un Portafolio Soberano

Pablo Tapia, Diego Vargas

DTBC – 1067

Macroeconomic Effects of Carbon-intensive Energy Price Changes: A Model

Comparison Matthias Burgert, Matthieu Darracq Pariès, Luigi Durand, Mario González, Romanos Priftis, Oke Röhe, Matthias Rottner, Edgar Silgado-Gómez, Nikolai Stähler, Janos Varga

DTBC - 1066

Bank Branches and the Allocation of Capital across Cities

Olivia Bordeu, Gustavo González, Marcos Sorá

DTBC – 1065

Effects of Tariffs on Chilean Exports

Lucas Bertinatto, Lissette Briones, Jorge Fornero

DTBC – 1064

Does Participation in Business Associations Affect Innovation?

Felipe Aguilar, Roberto Álvarez

DTBC – 1063

Characterizing Income Risk in Chile and the Role of Labor Market Flows

Mario Giarda, Ignacio Rojas, Sergio Salgado

DTBC – 1062

Natural Disasters and Slow Recoveries: New Evidence from Chile

Lissette Briones, Matías Solorza

DTBC – 1061

Strategic or Scarred? Disparities in College Enrollment and Dropout Response to Macroeconomic Conditions

Nadim Elayan-Balagué

DTBC – 1060

Quantifying Aggregate Impacts in the Presence of Spillovers

Dave Donaldson, Federico Huneus, Vincent Rollet

DTBC – 1059

Nowcasting Economic Activity with Microdata

Diego Vivanco Vargas, Camilo Levenier Barría, Lissette Briones Molina

DTBC – 1058

Artificial Intelligence Models for Nowcasting Economic Activity

Jennifer Peña, Katherine Jara, Fernando Sierra

