

What Drives Cross-Country Differences in the Share of Hand-to-Mouth Households? ^{*†}

Clara Arroyo[‡]

IMF

Esteban Tisnés[§]

Banco Central del Uruguay
Universidad ORT Uruguay

February 27, 2026

Abstract

Recent literature has highlighted that differences across countries in the share of Hand-to-Mouth households (HtM) are important for the transmission of aggregate shocks as well as monetary and fiscal policy. How can we explain cross-country differences in the share of hand-to-mouth households? In this paper, we first document significant heterogeneity in the share of HtM households across European countries. This heterogeneity is driven by large differences in the share of wealthy HtM households, who hold illiquid but no liquid wealth. On the contrary, the share of poor HtM, who hold neither liquid or illiquid wealth, is similar across countries. Second, we develop a two-asset life-cycle model with incomplete markets and uninsurable income risk and calibrate it to Spain. Through the lens of the model we study the role of country differences in income risk, the life-cycle profile of earnings, retirement benefits and financial conditions such as the spread between borrowing and savings rates and borrowing limits. Although we report substantial differences across countries in mean income, risk and retirement benefits, results suggest that they cannot explain cross-country differences in the share of HtM by themselves. On the other hand, differences in financial conditions are able to explain, *ceteris-paribus*, 51% of the differences in HtM shares, by explaining 68% of W-HtM shares.

*JEL Codes: E21, D14, D31

[†]We are incredibly grateful to Sebastián Fanelli for his support and guidance at all stages of this project. The paper has also benefited greatly from conversations with Josep Pijoan-Mas, Nezih Guner, Morten Ravn, Rody Manuelli, Francesco Chiocchio, and Luigi Maria Briglia. Clara Arroyo gratefully acknowledges funding from Santander's Research Chair at CEMFI. Esteban Tisnés acknowledges financial support from Grant PRE2021-097575 funded by MCIN/AEI/ 10.13039/501100011033 and by "ESF +", Fundación Carolina and CEMFI. The views expressed herein are those of the author and should not be attributed to the IMF, its Executive Board, or its management. The views expressed in this manuscript are those of the authors and do not necessarily represent the views of the Central Bank of Uruguay. All errors and omissions are ours only.

[‡]Contact: marroyo@imf.org

[§]Contact: etisnes@bcu.gub.uy

1 Introduction

The Marginal Propensity to Consume (MPCs) out of transitory income changes has become a central concept in modern macroeconomics. MPCs heterogeneity is related to the distribution of liquid assets: households who hold little liquidity, often referred to as Hand-to-Mouth (HtM), tend to have high MPCs¹. This can explain why countries where more households are classified as HtM experience larger aggregate responses to monetary policy shocks and have higher fiscal multipliers (??). The literature has focused mostly on understanding the consequences of countries having a high share of HtM households, rather than explaining why some countries exhibit higher shares of such households compared to others.

In this paper, we study quantitatively what drives cross-country differences in the share of HtM households. We first report substantial heterogeneity in the share of HtM households across European countries, particularly in the share of Wealthy HtM, those that hold illiquid assets but little liquidity. We then use a standard heterogeneous agent model with two assets to explore the determinants of these differences. We calibrate the model to Spain and perform different quantitative exercises. We find that differences in the life-cycle profile and transitory shocks that households face can explain only a moderate share of the differences observed in the data. On the other hand, differences in financial conditions have more potential to drive the cross-country heterogeneity in the share of HtM.

In the first part of the analysis, we estimate the share of HtM households for twenty-two European countries using the Eurosystem's Household Finance and Consumption Survey (HFCS). We find substantial heterogeneity across countries, consistent with previous literature. The total share of HtM can be as low as 10% in countries like Malta, Austria, or the Netherlands, and as high as 50% in countries like Hungary, Cyprus, Croatia, or Latvia. We further classify households into Poor (P-HtM) and Wealthy (W-HtM) depending on their holdings of illiquid assets, like real estate. While the P-HtM do not have either liquid

¹? perform a meta-analysis over empirical estimates of MPCs and find that the most robust result is the negative relationship between MPCs and liquidity or wealth.

or illiquid assets, the W-HtM can hold substantial amounts of illiquid assets but still have high MPCs because they hold little liquidity. We find that the cross-country difference in the total share of HtM is mostly driven by the W-HtM, as the P-HtM represent a similar proportion across countries. This highlights the relevance of following ? in looking at liquid and illiquid assets separately, instead of classifying agents as HtM based on their net worth.

We then analyze how the shares of W-HtM and P-HtM households correlate with some relevant country characteristics. As expected, we find that the share of W-HtM has a strong negative correlation with GDP per capita and financial development, proxied by an index produced by the International Monetary Fund (IMF). We look also at the correlation with some moments of the income distribution computed using data from the Survey on Income and Living Conditions (EU-SILC) conducted by the European Union. Both shares are positively correlated with the variance of income and income growth over the life cycle. Finally, we also find a negative cross-country correlation between expected replacement rates and the share of W-HtM. These results guide the selection of the most important features to include in the model, as well as the quantitative exercises we perform.

In the second part of the analysis, we quantitatively explore the role of different dimensions in shaping the share of HtM, as well as the decomposition between the wealthy and the poor. We develop a life-cycle model with incomplete markets, uninsurable income risk, and two assets, where the share of HtM households is determined endogenously. Households' earnings are composed of a deterministic age profile, a stochastic component with a transitory and a persistent shock, and a retirement scheme. Households can hold two types of assets with different returns and different degrees of liquidity. Borrowing is possible in the liquid asset up to a borrowing constraint and at a higher interest rate than the one that prevails for savings. These frictions give rise to an endogenous fraction of households that are HtM, meaning that they have zero or almost zero liquid assets and high MPCs.

We calibrate the model to Spain, a country that ranks in the middle of the HtM country-level distribution. We characterize the age profile of income and estimate the parameters

of the stochastic component of income using data from EU-SILC. We complement this with data from the OECD on replacement rates and minimum and maximum pensions to characterize the pension scheme. Finally, we calibrate internally another subset of parameters following a Simulated Method of Moments (SMM) approach, targeting moments related to liquid and illiquid asset holdings, particularly the share of wealthy and poor HtM.

Using the calibrated model, we perform several quantitative exercises to assess the importance of the different dimensions outlined above in explaining cross-country differences in the share of HtM. First, we explore the role of income dynamics. We estimate the same income process for all countries in the sample and find important differences in parameter estimates. We then compare the share of W-HtM and P-HtM in the baseline calibration for Spain against the counterfactual shares that arise when we impose the parameters of the income process estimated for the other countries. Overall, differences in the estimated income processes have limited power to explain cross-country differences in HtM shares. The share of W-HtM is most responsive to changes in the transitory component of risk, while the share of P-HtM responds more to changing the deterministic age-profile of income.

We then explore the role of financial conditions focusing on the spread between savings and borrowing interest rates, and the borrowing limit. Considering all countries in the sample, financial conditions by themselves are able to explain on average 51% of the differences in the shares of HtM. This is mainly through the W-HtM share (68% explained on average) with a smaller explanatory power in the case of the P-HtM (17%). Their explanatory power is stronger among countries with lower HtM shares, and gradually declines as the share increases. We show how other frictions, related to the housing and mortgage markets, may play a larger role for high HtM countries. We leave its quantification to further research.

Finally, we analyze whether differences in preferences across countries can drive the heterogeneity in HtM shares. Changing the discount rate has a strong effect on the share of W-HtM and P-HtM, making it a less appealing explanation as the share of P-HtM is quite

similar across countries. On the other hand, differences in initial conditions (distribution of assets at the beginning of working life) and demographic characteristics (death rates) have a minimal effect on HtM shares.

Our findings have relevant policy implications, in particular with respect to the regulations and policies related to households borrowing costs. Understanding why borrowing costs are higher in countries with higher shares of HtM, and taking actions to reduce those can have important macroeconomic stabilization effects. Through reducing the share of HtM, and the average marginal propensity to consume of the economy, countries can become more resilient to shocks and foster economic growth. As a consequence pro-competitive policies in the banking sector like open banking and the defense of the consumer of financial services, have important macroeconomic stabilization effects that are not usually taken into account.

Related literature. On the theory side, there is a burgeoning literature studying the effects of monetary and fiscal policy in the presence of HtM consumers in the New Keynesian framework. These models require a large share of HtM consumers to reproduce the large fiscal multipliers or responses to monetary policy observed in the data (e.g. ???). Empirically, ?, ? and ? find larger aggregate responses to fiscal and monetary policy shocks in countries where the share of HtM is larger. Both approaches take the share of HtM as exogenous, so we contribute to this literature by digging deeper into the drivers of HtM shares across countries.

This paper also relates to the literature studying determinants of consumption behavior for low-liquid wealth households. ? perform a meta-analysis on MPCs estimations and find that the most robust finding is that liquidity constraints imply larger MPCs. On the other hand, ? highlight the role of preference heterogeneity when trying to match a broader set of moments related to consumption spending. This literature has focused on consumers within specific countries, whereas we compare across countries. This way we contribute to this literature by revisiting the role of income dynamics, and other potential aggregate

drivers, in a cross-country setup.

This paper also contributes to the literature on cross-country variations in consumption and saving behavior. We empirically study saving behavior and liquidity constraints across European countries using the HFCS, finding that tax structures and social welfare systems significantly influence saving patterns. They find that Mediterranean countries, in particular, report higher liquidity constraints compared to Continental Europe. We study how household financial decisions shape the effects of monetary policy, highlighting how diverse country characteristics result in heterogeneous MPC distributions. Their findings indicate that Italy and Spain exhibit higher MPCs across income and education groups compared to France and Germany, attributing this to higher consumption floors and less educated households facing larger and more persistent permanent shocks.

2 The Hand-to-Mouth across Europe

In this section, we compute the fraction of Hand-to-Mouth households for a large sample of European countries² using data from the Eurosystem’s Household Finance and Consumption Survey (HFCS). We document substantial heterogeneity in the total share of HtM households and, particularly, in the share of the Wealthy HtM.

2.1 Data

To measure the shares of Hand-to-Mouth households in the data we rely on the Eurosystem’s Household Finance and Consumption Survey (HFCS). The survey is conducted by the national authorities of each country and is then harmonized across the European Union. The survey design follows the US’s Survey of Consumer Finances (SCF). There have been 4 waves conducted with data releases in 2013, 2016, 2020, and 2023, with ap-

²The countries in our sample are: Austria (AT), Belgium (BE), Croatia (HR), Cyprus (CY), Estonia (EE), Finland (FI), France (FR), Germany (DE), Greece (GR), Hungary (HU), Ireland (IE), Italy (IT), Latvia (LV), Lithuania (LT), Luxembourg (LU), Malta (MT), Netherlands (NL), Poland (PL), Portugal (PT), Slovakia (SK), Slovenia (SI), and Spain (ES).

proximately two years of gap between data collection and the release date. In what follows we present results using the wave released in 2020, the last wave before the COVID pandemic. We show in Figure 17 in the Appendix that the share of HtM and the differences across countries have been relatively stable over the different waves.

2.2 Identifying the Hand-to-Mouth Households

Agents are HtM if they spend all of their disposable income in every pay period, which means that they have a high marginal propensity to consume (MPC) out of transitory income changes. We classify HtM agents into poor and wealthy following ?, which has become standard in the literature. The poor HtM are agents with zero net worth and the wealthy HtM are those that hold sizable amounts of illiquid wealth, yet optimally choose to consume all of their disposable income during a pay period and carry zero liquid assets from one period to the next.

Following ? and ?, we measure the HtM in the data through the lens of a two-asset model. Households are HtM if they choose to be at one of the kinks in their budget constraint, either having zero liquid wealth ("zero kink") or being at the credit limit ("credit kink"). They are then further classified as wealthy if they hold positive illiquid wealth and poor if they hold zero illiquid wealth.

To identify households at the zero kink, we would ideally like to observe liquid balances at the exact moment before the household receives its monthly/weekly payment, as it reflects the amount of liquidity they choose to carry to the next period. However, our survey data does not allow this, as households' answers refer either to the average balance they hold in their accounts or to the balance at the time of the interview. This introduces some measurement error in the identification of HtM households. To tackle this issue we proceed as in ?, where a household is identified as HtM if its average liquid balances are below half of monthly income. This is assuming that households receive their income at the beginning of the month and spend it at a constant rate. Let b_i and a_i denote liquid and illiquid assets, y_i denote income, and \underline{b}_i be the credit limit for household i . A household is identified as

Table 1: Summary Statistics for Hand-to-Mouth across Europe

	Mean	Median	SD	Min	Max	N
Hand-to-mouth	0.31	0.27	0.15	0.10	0.65	22
Wealthy Hand-to-mouth	0.21	0.16	0.14	0.04	0.50	22
Poor Hand-to-mouth	0.10	0.10	0.04	0.04	0.17	22

Note: Own calculations based on our estimates of HtM status for households in the HFCS (third wave).

HtM if:

$$0 \leq b_i \leq \frac{y_i}{2} \quad (2.1)$$

or if

$$b_i \leq \frac{y_i}{2} - \underline{b}_i, \text{ and } b_i \leq 0 \quad (2.2)$$

A household is then classified as W-HtM if it additionally holds positive illiquid balances ($a_i > 0$) and as P-HtM otherwise.

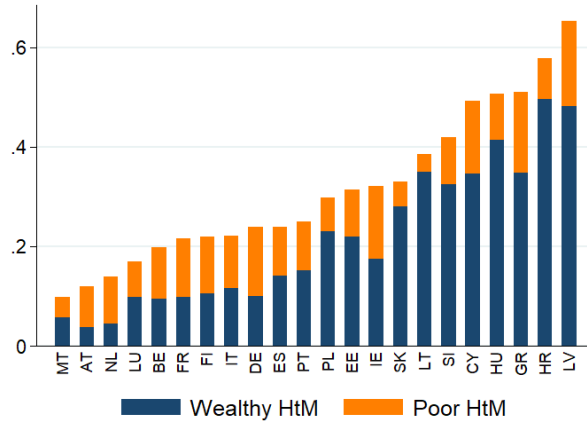
We follow ? in making some measurement assumptions. In the baseline measurement, we set the credit limit to one monthly income and the pay period to one month. We also follow their classification of assets and liabilities by liquidity. Liquid assets consist of balances held on sight and savings accounts, directly held mutual funds, publicly traded stocks and bonds (corporate and government), and cash³. Credit card debt and account overdrafts are considered liquid debt. Illiquid assets are mainly real estate net of mortgages and we include also retirement accounts, life insurance, and savings bonds. We do some robustness checks on these assumptions in Appendix A.3. Overall, we find that the differences in HtM shares across countries are maintained.

2.3 The Share of HtM Households across Europe

Table 1 reports some summary statistics for the share of HtM households across European countries estimated using the third wave. The average share of HtM households across

³Imputed following ? except for Spain.

Figure 1: The Share of HtM in Europe



Note: Based on our estimates of HtM status for households in the HFCS (third wave). Countries are: Austria (AT), Belgium (BE), Croatia (HR), Cyprus (CY), Estonia (EE), Finland (FI), France (FR), Germany (DE), Greece (GR), Hungary (HU), Ireland (IE), Italy (IT), Latvia (LV), Lithuania (LT), Luxembourg (LU), Malta (MT), Netherlands (NL), Poland (PL), Portugal (PT), Slovakia (SK), Slovenia (SI), and Spain (ES).

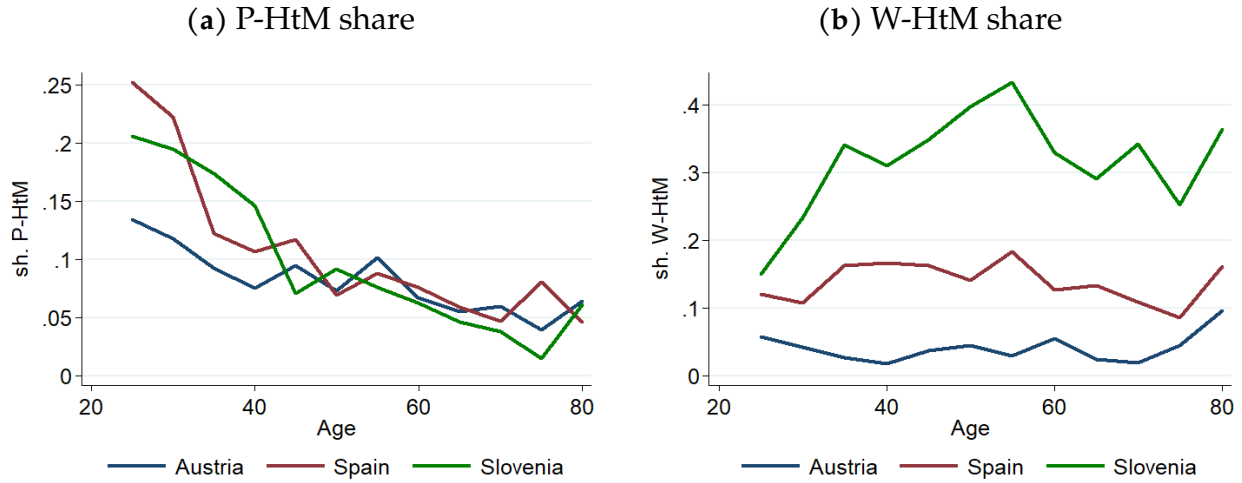
countries is 31%, with a standard deviation of 15%. The share of HtM ranges from 10% for Malta to 65% for Latvia. The values for all countries can be seen in Figure 1. Apart from Malta, other countries such as Austria and the Netherlands have low levels around 10%, and many countries have high values other than Latvia, above 50% for Hungary, Greece, and Croatia.

Table 1 also shows that the fraction of W-HtM is, on average, larger than the share of P-HtM (21% and 10% respectively). There is also more variability across countries in the share of W-HtM than in the share of P-HtM (14% and 4% standard deviation respectively). We can also see in Figure 1 how most of the difference in the share of HtM across countries is driven by the wealthy. For this reason, in the rest of this paper, we will put more emphasis on explaining the differences in the share of W-HtM rather than in the share of P-HtM.

2.4 Descriptive Statistics

We now show how the shares of W-HtM and P-HtM vary by age group in the different countries. To keep the exposition simple we select three countries: one with a low share of HtM (Austria), one with a high share of HtM (Slovenia), and one in the middle (Spain).

Figure 2: HtM Shares by Age



Note: This figure plots the share of poor hand-to-mouth (panel a) and wealthy hand-to-mouth (panel b) for different age bins for three selected countries (Austria, Spain, and Slovenia).

Figure 2 shows the share of W-HtM and P-HtM at different ages of the household head⁴. As shown in Panel (a), the pattern for the share of P-HtM is remarkably similar for the three countries: it decreases markedly with age. On the other hand, there are some differences in the pattern for the wealthy, shown in Panel (b). The share of W-HtM is mostly flat in Austria, while in Spain and Slovenia, it is increasing when young, peaking some years before retirement, and decreasing to a lower level afterward. In appendix B we also report some descriptive statistics related to households' education, marital status, balance sheets, and income over the life cycle.

2.5 Correlations with Country Characteristics

Ahead of the quantitative exercise in Section 5, we explore empirically here what country characteristics are associated with higher shares of HtM households. Table 11 in the Appendix reports the cross-country correlations of HtM, W-HtM, and P-HtM shares with the different characteristics explored here.

⁴The results shown here are obtained without controlling for age or year dummies. Doing so does not modify results significantly.

Figure 3: GDP and the share of HtM

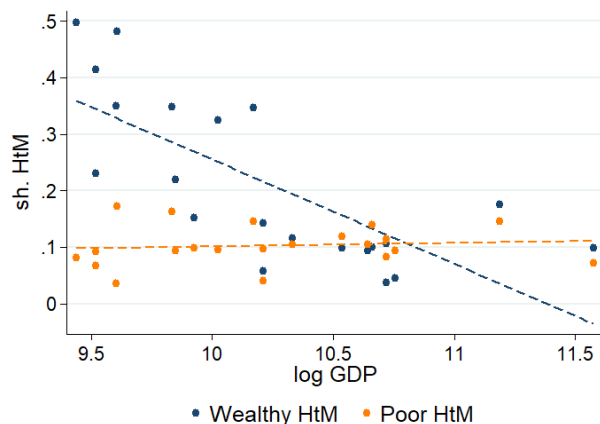
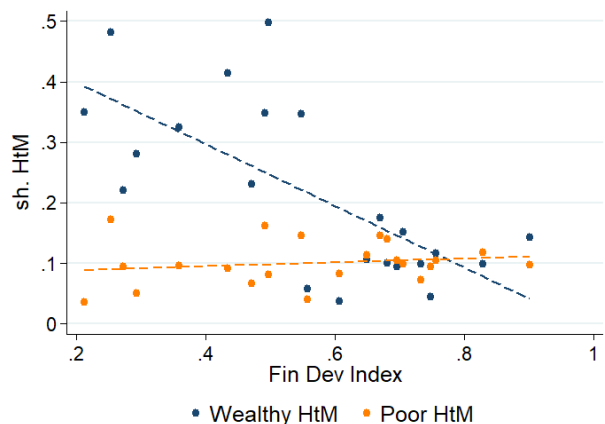


Figure 4: Financial Development and HtM

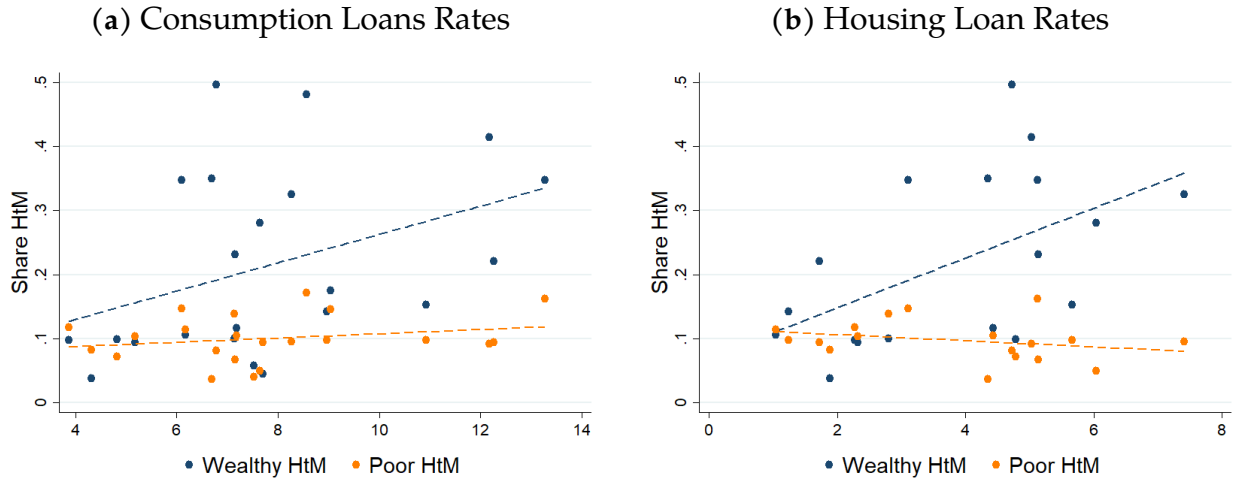


Note: This figure plots countries' share of wealthy (blue dots) and poor (orange dots) hand-to-mouth against their (log) GDP per capita (on panel a) and their financial development (on panel b), based on the financial development index produced by the IMF.

Aggregate Income. Figure 3 shows a clear negative relation between development (measured by the log of GDP per capita) and the share of HtM households. Richer countries tend to have a lower share of W-HtM households, while the share of P-HtM seems to correlate little with GDP. However, this correlation is noisy and there are significant differences in the share of HtM between countries with similar levels of aggregate income.

Financial Development. To analyze the correlation between HtM shares and financial development we use the financial development index produced by the IMF, which ranks countries according to the depth, access, and efficiency of their financial markets and institutions. The index uses data on bank credit to the private sector, pensions and mutual fund assets, access to ATMs, stock markets, and capital markets. Figure 4 shows that there is a negative association between financial development and the share of households that are W-HtM, while no correlation with the share of P-HtM. This points to a role of financial frictions in the determination of the share of W-HtM. Figure 5 also shows that higher borrowing rates (for consumption and housing loans) are associated with higher shares of W-HtM households at the country level.

Figure 5: HtM and Borrowing Rates



Income Dynamics. We now explore the relationship between HtM shares and different aspects of income dynamics at the country level. As can be seen in panel (a) of Figure 6, the variance of income is positively associated with both the share of W-HtM and P-HtM which indicates some role for income inequality behind the heterogeneity in HtM shares⁵.

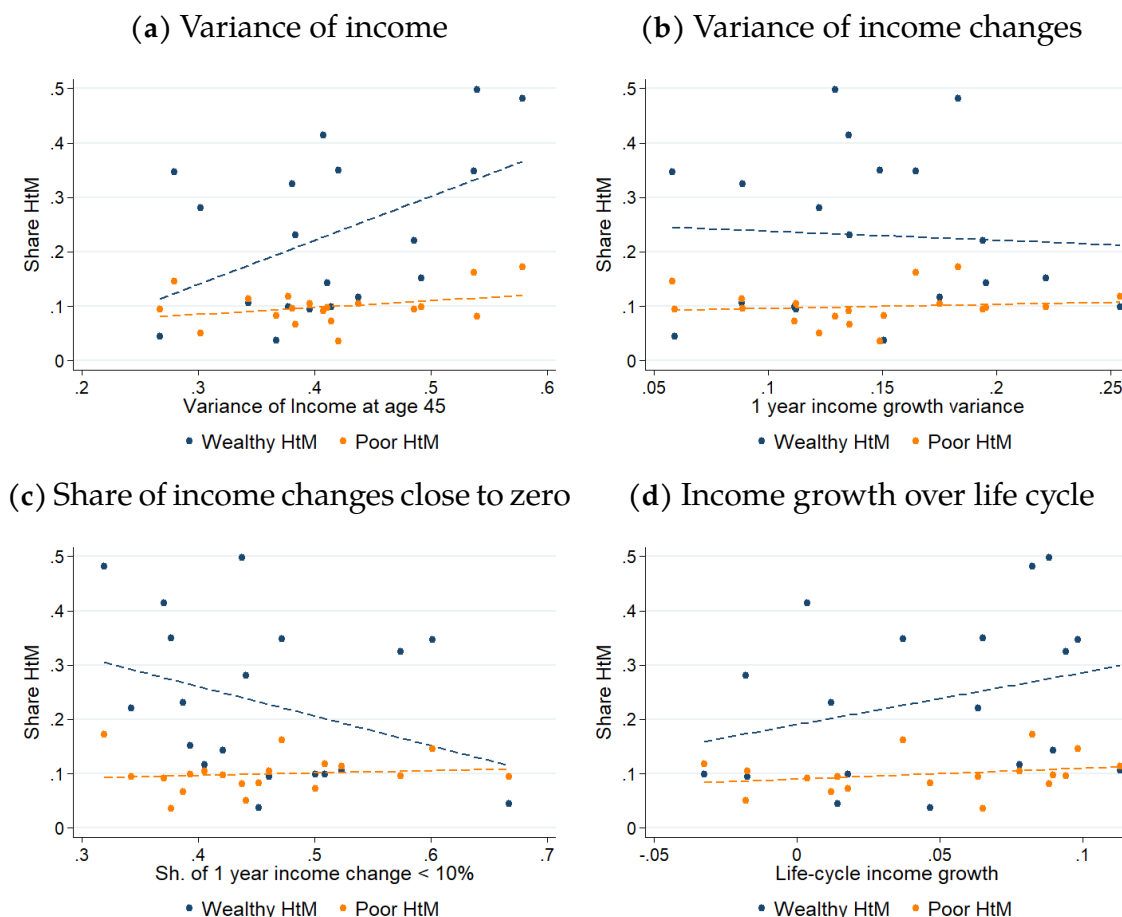
In panels (b) and (c) we show the correlation with two moments related to income risk. Panel (b) plots the share of HtM against the variance of one-year income changes where we see no clear correlation⁶. If anything, there is a negative correlation, consistent with higher income risk providing households an incentive to hold more liquid assets. Panel (c) plots the share of HtM against the share of one-year income changes that are below 10%, which is a measure of how concentrated income changes are around zero⁷. This moment can be informative about two different things with different implications for the share of HtM. First, a high share of small income changes can mean that income risk is low and households have little incentive to hold liquid assets. This would point to a positive correlation with the share of HtM, unlike what we see in the data. Second, a high share

⁵Here we use the variance of income at age 45 but the picture is similar if we look at other ages. Figure 24 in the Appendix reports the same figures for the variance of income at ages 25 and 55

⁶In this figure we use the variance of one-year changes in income. However, we see a similar lack of correlation if we look instead at two-year and three-year changes, as can be seen in Figure 25 in the Appendix.

⁷In Figure 26 of the Appendix we use other thresholds (changes less than 20% and less than 50%) and we obtain very similar results.

Figure 6: Income Dynamics and the Share of HtM

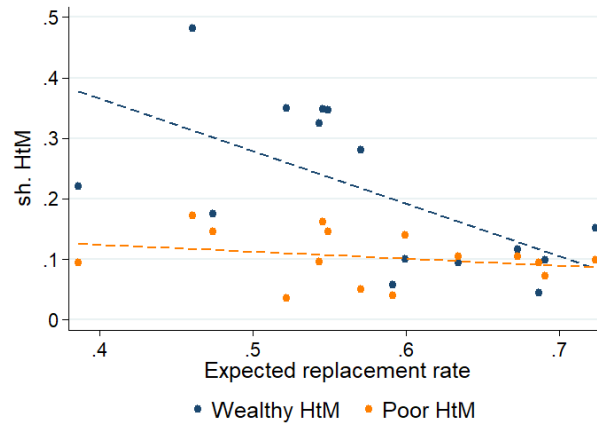


Note: This figure plots countries' share of wealthy (blue dots) and poor (orange dots) hand-to-mouth against their some moments related to income dynamics: the variance of income at 45 years old (panel a), the variance of income changes over one year (panel b), the share of income changes below 10% (panel c) and income growth between 25-30 and 50-55 years old (panel d).

could mean that most of the shocks that households face are small rather than large. Since liquid assets are better for insuring against small frequent shocks and illiquid assets are better for insuring against large infrequent shocks, this would favor holding liquid over illiquid assets. This would point to a negative correlation with W-HtM shares, as we see in the data. In short, it is hard to separate the size of shocks from the relative frequency of small and large shocks in the data, so we will explore this further in the quantitative section.

Finally, we turn to the life-cycle profile of income. Panel (d) shows the correlation of HtM

Figure 7: Expected replacement rates and the share of HtM



Note: This figure plots countries' share of wealthy (blue dots) and poor (orange dots) hand-to-mouth against the average expected replacement rate, computed from the HFCS data.

shares with income growth between ages 25-30 and 50-55. A steeper life-cycle profile should give households an incentive to borrow at younger ages to smooth consumption over the life-cycle, pushing some households to the borrowing limit. This would then be consistent with a higher share of HtM households, which we see in the data for both types of HtM, stronger in the case of the W-HtM.

Retirement. Finally, we explore the relationship between the share of HtM households and retirement conditions. Retirement is a large one-time event on which households have a lot of information. Households can make reasonably good forecasts of their income after retirement, especially when they are approaching the age of retiring, influencing their consumption and saving decisions.

Retirement can also provide different incentives across the income distribution. Income-rich households usually experience a drop in their income when they retire, meaning that their replacement rate is well below one. This gives them incentives to save in high return assets to smooth consumption before and after retirement, increasing the likelihood that they become W-HtM. On the other hand, income-poor households have larger replacement rates, which can even be larger than one if minimum pensions are large enough.

As a consequence they have an incentive to borrow against their future higher income, increasing the likelihood that they become P-HtM. Using data from the HFCS on households' expected replacement rate, we show in Figure 7 that there is a negative cross-country correlation between expected replacement rates and the share of W-HtM⁸.

3 Model

We build a life-cycle model with incomplete markets and idiosyncratic income risk in the tradition of ?, ?, and ?. We focus on the consumption-saving problem of a household that has access to two assets with different degrees of liquidity as in ?. During their working life, they face uninsurable idiosyncratic risk in the form of transitory and persistent income shocks. When households retire they receive a pension that is determined by their average income during their working life. Households face early death risk during their whole lifetime but this risk becomes more relevant after retirement when it is the only risk they face.

3.1 Household's Problem

Households maximize their expected utility at age t . If they survive (with probability $1 - \delta_t$) they receive utility from consuming c_t and if they die (with probability δ_t) they receive utility from leaving liquid (b_t) and illiquid (a_t) bequests. Preferences are time-separable and future utility flows are discounted at rate ρ . The household lives from age $t = 0$ to age $t = T$ and maximizes:

$$\mathbb{E}_0 \int_0^T e^{-\rho t} \left((1 - \delta_t)u(c_t) + \delta_t \tilde{u}(a_t - \underline{a}, b_t) \right) dt \quad (3.1)$$

⁸Table 13 in the Appendix shows the results of regressing HtM status on expected replacement rates at the household level. We see that higher expected replacement rates are associated with a lower probability of being HtM, particularly at later ages.

where

$$u(c_t) = \frac{1}{1-\gamma} c_t^{1-\gamma} \quad (3.2)$$

$$\tilde{u}(a_t + \underline{a}, b_t) = \kappa_a \log(a_t + \underline{a}) + \kappa_b \log(b_t) + \kappa_{b^{pos}} \mathbb{1}\{b_t > \bar{y}\} + \kappa_{nw} \mathbb{1}\{b_t + a_t < 0\} \quad (3.3)$$

The utility function $u(c_t)$ is Constant Relative Risk Aversion (CRRA) with parameter γ and the utility of leaving bequests $\tilde{u}(a_t - \underline{a}, b_t)$ is outlined in equation 3.3. The parameters κ_a , κ_b , and $\kappa_{b^{pos}}$ weight the utility of leaving bequests in each asset. Following ?, \underline{a} introduces a non-homotheticity in illiquid bequests so that only wealthy households receive utility from leaving illiquid assets as bequests.

The utility of leaving liquid bequests has two components: a standard warm glow motive and an extra utility of leaving liquid assets above some positive threshold. This gives households an extra reason to demand liquidity at older ages, which we use as a reduced form for other reasons to hold liquidity⁹. Finally, liquid bequests are allowed to be negative as long as net worth remains positive. We incorporate this by including a penalty ($\kappa_{\text{net worth} < 0}$) whenever net worth is negative.

Households maximize their utility subject to a budget constraint, a credit limit, a no short-sale constraint on the illiquid asset, and the transaction cost of accessing the illiquid account. Assets holdings of the household evolve according to:

$$\dot{b}_t = Z_t + r^b(b_t) b_t - d_t - \chi(d_t, a_t) - c_t \quad (3.4)$$

$$\dot{a}_t = r^a a_t + d_t \quad (3.5)$$

$$a_t \geq 0, \quad b_t \geq \underline{b} \quad (3.6)$$

⁹After retirement households in the model no longer have income risk and face only death risk. In reality, other risks also shape the patterns of wealth accumulation in old age, like health shocks and medical expenditures as in ?. We can think of these as being included in the second component of the liquid bequests, whose parameters we will calibrate to match the demand for liquid assets at old age and the share of W-HtM in the last age bin of our sample.

Households take interest rates as given, where there is a spread between saving and borrowing rates in the liquid asset ($r^b(b < 0) = r^b(b > 0) + \kappa$). Z_t denotes income, d_t are deposits to the illiquid account, and $\chi(d, a)$ are transaction costs associated with a given level of deposits and illiquid account balances. As in ? we adopt the following functional form:

$$\chi(d, a) = \chi_0|d| + \chi_1 \left(\frac{d}{a}\right)^{\chi_2} a \quad (3.7)$$

The transaction cost has two components: a linear and a convex component. The linear component (with $\chi_0 > 0$) creates an inaction region, where the cost of depositing/withdrawing is larger than the marginal gain of depositing/withdrawing so households choose $d = 0$. On the other hand, the convex component ensures that deposits are finite.

During their working life, households' income is given by two components, a deterministic part that only depends on age ($\mu_z(t)$) and a stochastic transitory-persistent income process (z_t). After retirement, households receive a pension that is determined by the average income in the last T^{pen} years of working, bounded by a minimum and a maximum pension.

Income is then:

$$Z_t = \begin{cases} z_t \mu_z(t) & \text{if } t \leq T^{ret} \\ \min\{\max\{\tilde{z}_{T^{ret}}, \tilde{z}_{min}\}, \tilde{z}_{max}\} & \text{if } t > T^{ret} \end{cases}$$

where $\tilde{z}_{T^{ret}}$ is computed according to:

$$\tilde{z}_t = \frac{1}{T^{pen}} \int_0^{t-(T^{ret}-T^{pen})} \mu_z(t-s) z_{t-s} ds$$

$$d\tilde{z}_t = \frac{1}{T^{pen}} \mu_z(t) z_t dt$$

As can be seen below, \tilde{z}_t is included as a state variable in the recursive formulation for the household problem, where households keep track of their "contributions" to their future

pension. This recursive formulation is:

$$\begin{aligned} \rho V(a, b, z_t, \tilde{z}, t) = \max_{c, d} & \left\{ U(c) + V_b (Z_t + r^b (b) b - d - \chi (d, a) - c) \right. \\ & \left. + V_a (r^a a_t + d_t) \right\} \\ & + \sum_{z'} \lambda(z, z') (V(a, b, z', \tilde{z}, t) - V(a, b, z, \tilde{z}, t)) \\ & + \frac{\mu_z(t) z_t}{T^{pen}} V_{\tilde{z}}(\cdot) + V_t(\cdot) \end{aligned}$$

The first-order conditions of this problem are given by:

$$\begin{aligned} U'(c) &= V_b \\ V_b (1 + \chi_d(d, a)) &= V_a \end{aligned}$$

We can obtain optimal deposits from the last equation:

$$\frac{d}{a} = \left\{ \left[\left(\frac{V_a}{V_b} - 1 + \chi_0 \right)^- \right]^{\frac{1}{x_2 - 1}} + \left[\left(\frac{V_a}{V_b} - 1 - \chi_0 \right)^+ \right]^{\frac{1}{x_2 - 1}} \right\} \chi_1 \chi_2^{\frac{1}{1 - x_2}} \quad (3.8)$$

Here it is clear that the larger χ_0 , the larger is the region where households choose not to deposit (or withdraw). As commented before, $\chi_0 > 0$ generates a no-adjustment region where the household does not withdraw or deposit because it is too costly to do so.

3.2 The Hand-to-Mouth in the Model

The model can generate HtM households through three frictions: the borrowing spread on the liquid asset, the transaction cost of accessing the illiquid account, and the borrowing limit. Households with zero liquid assets are at a kink in their Euler equation that makes consumption smoothing not optimal in some cases. For example, a high borrowing rate can make smoothing too expensive for P-HtM households when they face a small negative

shock. Additionally, for W-HtM households who could liquidate some of their illiquid assets to consume more, the adjustment cost may be too high. A similar argument holds when W-HtM households receive a small positive shock. Finally, households at the credit limit are simply not able to increase their debt to smooth consumption when hit by a negative shock. In all these cases, the consumption of HtM households will closely follow their income changes, and they exhibit high MPCs.

3.3 A Transitory-Persistent Income Process

We describe here in more detail the stochastic component of income, which is key in the model because earnings risk is the only reason to hold liquidity. We model the continuous-time earnings dynamics in the model following z_{it} , characterized by transitory and persistent shocks with different frequency, persistence, and size. What matters for the portfolio decision of households is not just the amount of risk (or the size of the shocks) that they face but also the type of shock. As we discussed before, liquid assets are better at insuring against small, frequent, and transitory shocks. On the other hand, illiquid assets are better at insuring against large, infrequent, and persistent shocks, since their higher return can compensate for the transaction cost that is paid whenever there is a large and/or persistent shock.

Log-earnings are the sum of two “jump-drift” processes ($z_{1,it}$ and $z_{2,it}$):

$$\log(z_{it}) = z_{1,it} + z_{2,it} \quad (3.9)$$

$$dz_{j,it} = -\beta_j z_{j,it} dt + dJ_{j,it} \quad j = 1, 2 \quad (3.10)$$

where jumps ($dJ_{j,it}$) arrive at a Poisson rate λ_j . Conditional on a jump, the new log-income level is drawn from a normal distribution with zero mean and variance σ_j^2 . When there are no jumps the process reverts to the mean at rate β_j .

4 Taking the Model to the Data

In this section, we describe how we take the model to the data. Our baseline calibration is for Spain, a country that is in the middle of the HtM distribution among the set of countries in our sample. We estimate the parameters of the exogenous income process using data from the Survey on Income and Living Conditions. Then we take a subset of parameters of the model from external sources and internally calibrate the remaining parameters to match some features of the data for Spain, mainly moments related to households' asset holdings and HtM status.

4.1 Income Process Estimation

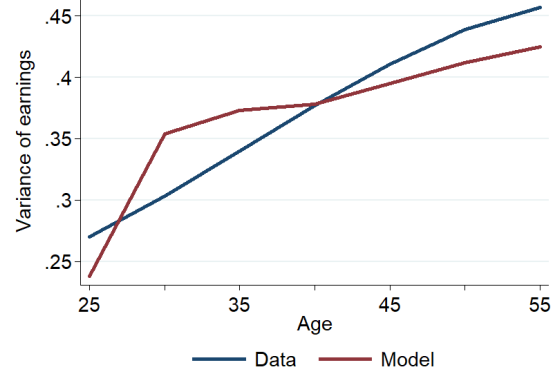
Data. We use data from the Survey on Income and Living Conditions conducted by the European Union. This survey has a panel dimension, where households are re-interviewed annually for 4 consecutive waves. The data cleaning consists of dropping households that are not present with positive income in all 4 waves and those whose head's age increases by more than two years from wave to wave. Additionally, we drop those on the top of the income growth distribution when computing income growth variances. Since we are interested in income at the household level, we define income as the sum of personal income of household members (head and spouse) and other income generated by the household (i.e. family or children-related allowances). Personal income comprises wages and other work-related payments, self-employment income, unemployment payments, and other social benefits.

Estimation. Once again our estimation strategy follows ?. We estimate the income process in equations (3.9)-(3.10) by Simulated Method of Moments (SMM). We target moments related to the distribution of income levels and income changes, which can be seen in Table 2¹⁰. As in ?, we seek to identify the frequency of shocks from low-frequency data

¹⁰Before computing the moments we residualize our measure of income with observable variables like education of the head and cohort effects.

Table 2: Income Process Estimation Fit

	Data	Model
$\text{Var}(y_{25})$	0.27	0.24
$\text{Var}(y_{30})$	0.3	0.35
$\text{Var}(y_{45})$	0.41	0.39
$\text{Var}(y_{55})$	0.46	0.43
1 year autocorr.	0.92	0.88
2 year autocorr.	0.79	0.8
3 year autocorr.	0.72	0.78
$\text{Var}(\Delta_1 y)$	0.2	0.2
$\text{Var}(\Delta_2 y)$	0.3	0.31
$\text{Var}(\Delta_3 y)$	0.35	0.34
$\text{Kurt}(\Delta_1 y)$	10.09	10.36
$\text{Kurt}(\Delta_3 y)$	7.54	6.7
Frac. $\Delta_1 y < 10\%$	0.42	0.52
Frac. $\Delta_1 y < 20\%$	0.6	0.64

Figure 8: Income Variance by Age

Note: Table 2 compares the targeted moments of the income distribution in the data and in the model for Spain. Figure 8 plots the variance of earnings for different age groups in the data and in the model, also for Spain.

(annual income). To do so we rely on high-order moments (kurtosis) of the income growth distribution at different horizons, such as changes over one and three years.

However, our data differs from the data used by ? in several dimensions. While they used social security data for the US from 2002 to 2006, we do not have access to this data for all countries in the European Union so we resort to survey data. Our data allows us to model directly household level earnings which is more suited to analyze wealth accumulation decisions at the household level. In particular, analyzing household-level earnings allows us to include natural insurance arrangements that arise in households with two earners. On the downside, the estimation of the moments will be more exposed to measurement error, especially for higher-order moments. For this reason, we assign less weight to these moments in the estimation algorithm.

Estimated parameters and Model fit. The values of the estimated parameters are presented in Table 3. The results show a clear differentiation between the two processes. One is characterized by frequent but small transitory shocks ($j=1$), whereas the other is charac-

Table 3: Income Process Estimated Parameters

	Parameter	$j = 1$	$j = 2$
Arrival rate	λ_j	0.324	0.002
Standard deviation	σ_j	1.018	1.371
Mean reversion	β_j	0.884	0.0004

terized by infrequent but large persistent shocks ($j=2$). This model for the income process fits the data quite well, as can be seen in Table 2, which reports the targeted moments in the model and the data. Figure 8 also compares the variance of income in the model and in the data for different ages (only ages 25, 30, 45, and 55 are targeted)¹¹.

4.2 Calibration Strategy

Table 4 shows the values of all the parameters apart from those of the income process outlined before. A first subset of parameters is chosen based on the literature and external sources. We set the inverse of the intertemporal elasticity of substitution to one and the liquid interest rate on savings to 0.5% annual. We set the borrowing limit to one average quarterly income and compute the survival probabilities (δ_t) using mortality data from the Human Mortality Database (?). Finally, we set the retirement age at 65 and have households live until the age of 90.

The remaining parameters are calibrated internally following an SMM approach and using data from the HFCS and Eurostat. As can be seen in Table 5, we target the mean and median of the wealth-to-income ratios for liquid and illiquid assets, the share of debtors, and the share of W-HtM and P-HtM households, including the decomposition between those at the zero kink and at the credit limit. We also target some moments for the last age bin (eighty years old) to inform the model about bequests. We again target the wealth-to-

¹¹In Spain and most countries the variance of income is increasing with age, consistent with income processes having high persistence. However, when we estimate the income process for all countries in our sample to perform the quantitative exercise featured in the next section we find that for some of them, the variance of income decreases in young ages and increases afterward (see Figure 27 in the Appendix). This is potentially related to education and late entry into the labor market. This feature biases the estimation of the model, and as a consequence, we model the first transitory and the first persistent shock with a different standard deviation to match the initial level of income variance (at age 25).

Table 4: Model Calibrated Parameters

Parameter	Description	Value	Calibration
<i>Preferences</i>			
γ	Inv. Intertemp. Elast. Subst.	1	External
ρ	Discount rate (p.a.)	6.4%	Internal
\underline{a}	Non-homothetic bequest	1000	Internal
κ_a	Illiquid bequest weight	3	Internal
κ_b	Liquid bequest weight	2.29	Internal
$\kappa_{b^{pos}}$	Positive liquid beq. utility	0.89	Internal
<i>Rates</i>			
r^a	Illiquid rate (p.a.)	7%	Internal
$r^b(b > 0)$	Savings rate (p.a.)	0.5%	External
$r^b(b < 0)$	Borrowing rate (p.a.)	6.3%	Internal
\underline{b}	Borrowing limit	-8000	External
<i>Adjustment cost function</i>			
χ_0	Linear component	0.03	Internal
χ_1	Convex component	0.97	Internal
χ_2	Convex component	1.73	Internal

income ratios, the shares of W-HtM and P-HtM and their decomposition, and additionally, we target the share of households leaving positive bequests in each asset, and the 30th percentile of illiquid assets bequests.

Table 5 compares the moments computed in the model to those computed with the data. The model matches the share of W-HtM and P-HtM, and the decomposition between those at the zero kink and at the credit limit. It also matches well the level of total wealth, although with more liquid wealth and less illiquid wealth than in the data. In terms of the moments computed for the last age bin, the model matches well the shares of W-HtM and P-HtM, as well as the fraction of households with positive liquid wealth. However, the model overstates how concentrated illiquid assets are at old age: it falls short of matching the share of households with positive illiquid wealth¹² and overstates the accumulation of

¹²Illiquid wealth consists mainly of housing, which we do not explicitly model. Usually to be able to match ownership rates (what we call fraction with positive illiquid balances) the literature resorts to the additional utility of owning houses, which we avoid in this paper.

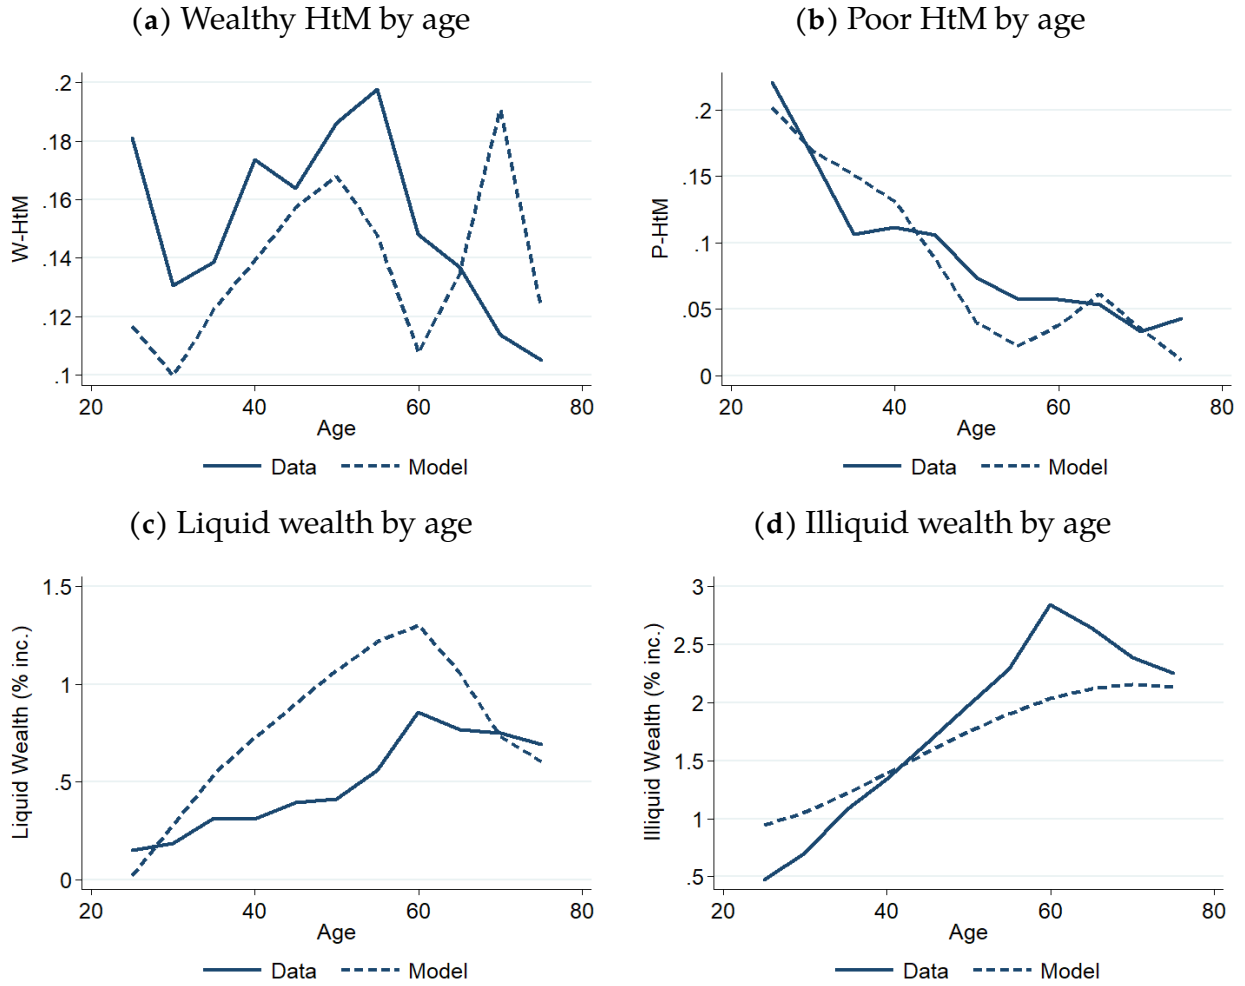
Table 5: Model Fit

Moment	Data	Model	Source
Targeted	<i>Average across ages</i>		
Liquid wealth (mean)	0.49	0.79	Eurostat
Illiquid wealth (mean)	1.79	1.66	Eurostat
Debtors (%)	0.31	0.27	HFCS
Total W-HtM (%)	0.16	0.14	HFCS
Zero kink (%)	0.14	0.12	HFCS
Credit limit (%)	0.02	0.02	HFCS
Total P-HtM (%)	0.09	0.09	HFCS
Zero kink (%)	0.08	0.08	HFCS
Credit limit (%)	0.01	0.01	HFCS
	<i>Last age bin</i>		
30th percentile of $f_{a a>0}$	0.95	1.85	Eurostat and HFCS
Share $a > 0$	0.91	0.65	Eurostat
Share $b > 0$	0.91	0.96	HFCS
Total W-HtM (%)	0.10	0.12	HFCS
Zero kink (%)	0.10	0.12	HFCS
Credit limit (%)	0.00	0.00	HFCS
Total P-HtM (%)	0.04	0.01	HFCS
Zero kink (%)	0.04	0.01	HFCS
Credit limit (%)	0.00	0.00	HFCS
Non-targeted	<i>Average across ages</i>		
Liquid wealth (median)	0.09	0.30	Eurostat
Illiquid wealth (median)	1.15	1.16	Eurostat

wealth at the top (the 30th percentile of the illiquid assets is twice as much as in the data).

Finally, we look at some non-targeted moments. The last panel of Table 5 shows that the model matches the median of illiquid wealth and overestimates the median of liquid wealth. Figure 9 illustrates the model fit in terms of the age profile of HtM status and wealth accumulation, which were not targeted in the calibration. The overall fit is good: the model captures the hump shape in the share of W-HtM, liquid wealth, and illiquid wealth, and the size and monotonic decrease in the share of P-HtM. However, the model features more liquid wealth than in the data until retirement. This is also reflected in a

Figure 9: Non-targeted Moments

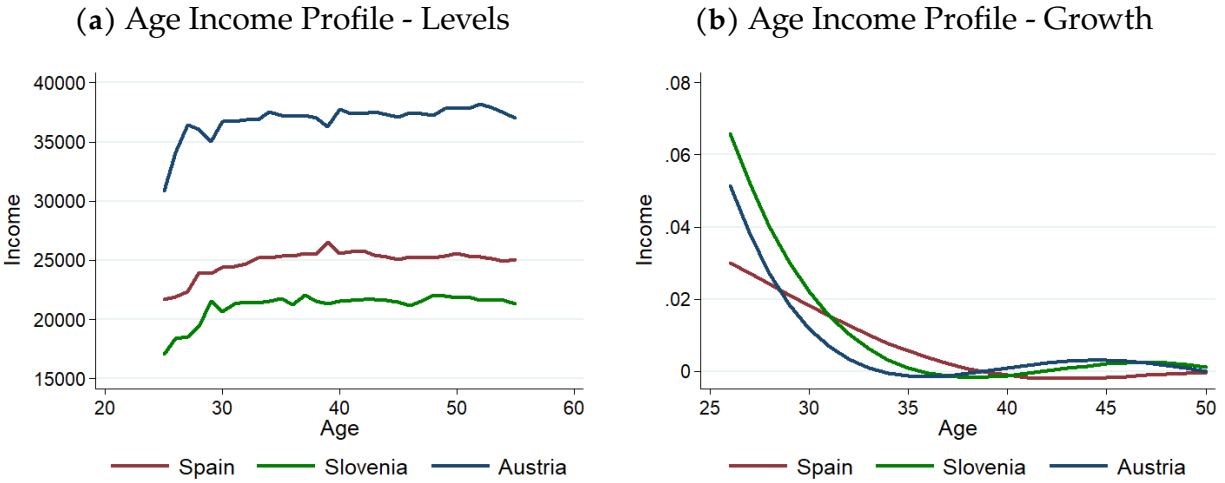


lower share of W-HtM at young ages in the model relative to the data. The share of W-HtM spikes immediately after retirement, when households no longer face income risk but only early death risk. Since the probability of death is not very high, this pushes down the demand for liquidity.

5 Quantitative Exercise

In this section, we use the calibrated model to explore the determinants of countries' share of HtM households. We first analyze the role of differences in income dynamics and then

Figure 10: Age profile of income



Note: Panel (a) plots average income for different age groups and for three countries (Spain, Slovenia, and Austria). Panel (b) plots the average income growth for the same age groups and countries.

explore other possible determinants, such as financial conditions, preferences, and demographic characteristics.

5.1 Income Process

Starting from the baseline calibration for Spain, we explore whether changing different components of income to those estimated for the other countries in the sample can contribute to explaining the differences in the share of HtM households. There are four dimensions of income in the model that can affect the share of HtM: the deterministic age profile of earnings, the transitory shocks, the persistent shocks, and the characteristics of the pensions system. We first describe the differences across countries in these dimensions and then report the results of the quantitative exercise.

Differences in the Income Process across countries. We focus on three countries to portray how the components of income can differ across countries and leave a complete set of comparisons to appendix C. As in section 2.4, we compare Spain (our baseline calibration), with Austria (with fewer HtM households) and Slovenia (with more HtM households).

Table 6: Income Process: Estimated Parameters

Shock	Parameter		AT	ES	SI
Transitory	Arrival Rate	λ_1	0.355	0.324	0.21
Transitory	Standard dev.	σ_1	0.862	1.018	0.763
Transitory	Mean Reversion	β_1	0.897	0.884	0.451
Persistent	Arrival Rate	λ_2	0.008	0.002	0.008
Persistent	Standard dev.	σ_2	1.635	1.371	1.34
Persistent	Mean Reversion	β_2	0.045	0.0004	0.015

We first compare these countries in terms of the life-cycle profile of income, which we include in the model as a deterministic age profile. In Figure 10 we compare the level of income, the age profile, and the (smoothed) growth rate. Households in richer countries have a more stable life-cycle profile of income on average. The model predicts that households with lower income save more to escape from the borrowing limit. Households who face a steeper income path, like in Slovenia, have an incentive to borrow to smooth consumption over the life cycle. As households can only borrow in liquid assets, this affects demand for liquidity and the likelihood of falling into HtM status.

Next, we compare countries in terms of the characteristics of the transitory and persistent shocks. Table 6 presents a comparison of the estimated parameters for the three countries. Transitory shocks are smaller (lower standard deviation), less frequent, and more persistent in Slovenia than in Spain, giving households fewer incentives to hold liquidity¹³. On the other hand, Austria has smaller transitory shocks that are slightly more frequent than in Spain, with an ambiguous effect on liquidity. It also has larger persistent shocks that are however less persistent than in Spain, with ambiguous implications for the holding of illiquid assets¹⁴.

Finally, we look at the pension system that determines income after retirement. There

¹³In Figure 29 in the appendix we compare the share of HtM for different values of λ_1 , σ_1 , and β_1 , showing that a lower arrival rate, lower standard deviation, and higher mean reversion of the transitory component are associated in the model with a higher share of W-HtM and P-HtM.

¹⁴In Figure 29 in the appendix we compare the share of HtM for different values of λ_2 , σ_2 , and β_2 , showing that changing these parameters has little effect on the share of HtM. If anything, a lower arrival rate is associated with a higher share of W-HtM and P-HtM, and a higher standard deviation and higher mean reversion are associated with a higher share of W-HtM and a lower share of P-HtM.

is a variety of ways pensions are funded across Europe. We focus on the compulsory contribution-based system as we want to capture retirement income that households do not need to actively save for. Households can choose to save for retirement to complement the state-provided pension. The incentives to do so will arise from their life-cycle income, especially if their income when they are close to retirement is larger than what the pension system provides. This incentive is partially captured by the maximum pension they will receive. On the other end of the distribution, households with low-income levels benefit from generous minimum pensions, which disincentivizes saving for retirement.

Table 7 shows the differences in some parameters of the pensions systems. Minimum pensions as a percentage of average earnings are around 22% in the whole sample, with Eastern European countries typically below this number (at around 10%-15%), and Central Europeans above it (around 30%). When present, maximum pensions are usually high enough that they are only binding for a small share of households. As a consequence of generous minimum pensions, replacement rates¹⁵ are high for low earners and decrease with income. This can be seen in the table, where we compare the replacement rate for earners with the average income to those for low-wage earners (at half the average income) and high-wage earners (at twice the average income). Replacement rates are below 100% because of the increasing age profile of income and maximum contributions. Since the model's endogenous replacement rates are above the one in the data, to match the average replacement rate we adjust the mean of the deterministic component of income during retirement.

Contribution to differences in HtM shares. We now report the results of the quantitative exercise, where we change the parameters of the income process from those of Spain (the baseline calibration) to those estimated for each of the other countries. Then we see how the fraction of HtM changes endogenously in the model.

In Figures 11 and 12, each dot represents the difference in the share of HtM households

¹⁵These are computed as pension income over the last years of earnings.

Table 7: Pensions systems

Variable	Reference	AT	ES	SI
Minimum	% of avg. earnings	28%	36%	33%
Maximum	% of avg. earnings	155%	181%	325%
Replacement rate	0.5x avg. Income	69%	74%	79%
Replacement rate	1x avg. Income	71%	74%	58%
Replacement rate	2x avg. Income	55%	69%	54%

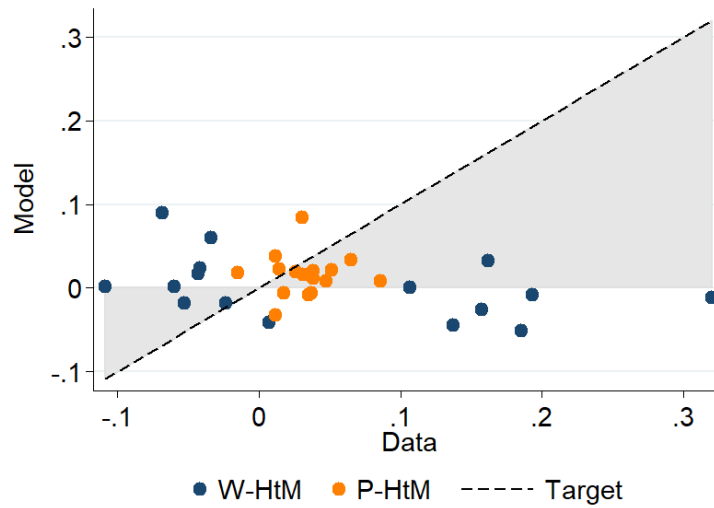
Note: This Table reports minimum and maximum pensions as a fraction of average countries for Austria, Spain, and Slovenia. The table also reports the average replacement rate for earners at half the average income, earners at the average income, and earners at twice the average income. Data is from ?.

(wealthy in blue and poor in orange) between a given country and Spain (baseline). In the horizontal axis is the difference observed in the data. In the vertical axis is the difference we obtain in the model when we change the parameters of the income process to those estimated for the corresponding country (leaving all other parameters unchanged). If dots fall in the 45-degree line then, only by changing the parameters of the income process, the model perfectly matches the difference between countries observed in the data. Instead, if dots fall in the horizontal axis, changing the income process has no effect on the share of HtM households. Finally, if dots are in the shaded area (between the 45-degree line and the horizontal axis), it means that the difference in the income process parameters can explain some, but not all, the difference between countries' HtM shares.

First, in Figure 11 we show the result of switching all components of the income process at the same time. Most of the blue dots lie outside the shaded area, meaning that, for most countries, differences in the income process cannot help explain differences in the share of W-HtM households. On the other hand, since many of the orange dots do fall in this area, the income process seems to explain more of the differences in the share of P-HtM.

Figure 12 shows the results of changing components of the income process one at a time. The most relevant component to explain the difference in the share of P-HtM is the life-cycle profile of earnings. For most countries, it contributes to bringing the share of P-HtM closer to the data, while pushing the share of W-HtM further away. The differences in the life-cycle profile can be summarized in differences in the overall level of income and

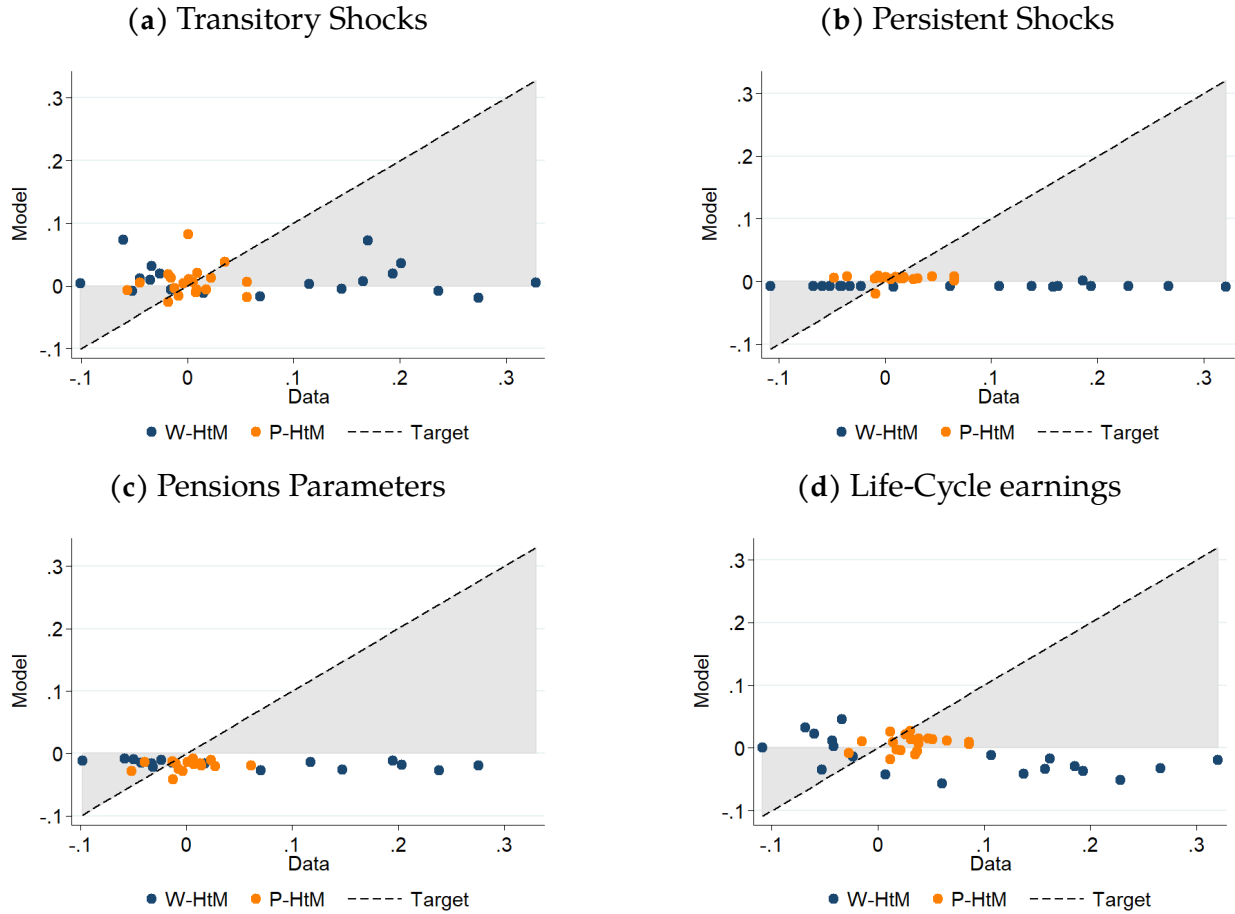
Figure 11: Switching all Components of Income



Notes: Each dot in the figure represents the difference in the share of HtM households (wealthy in blue and poor in orange) between a given country and Spain (baseline). In the horizontal axis is the difference observed in the data. In the vertical axis is the difference we obtain in the model when we change all the parameters of the income process to those estimated for the corresponding country (leaving all other parameters unchanged). The shaded area between the horizontal axis and the 45-degree line represents the area where changing only the corresponding parameters gets the model closer to reproducing the difference observed in the data between Spain and the corresponding country.

the slope over age. If the earnings profile is flatter, there is a weaker incentive to borrow against future income, which increases liquidity. On the other hand, a poorer overall level of income should increase precautionary savings to escape the borrowing constraint.

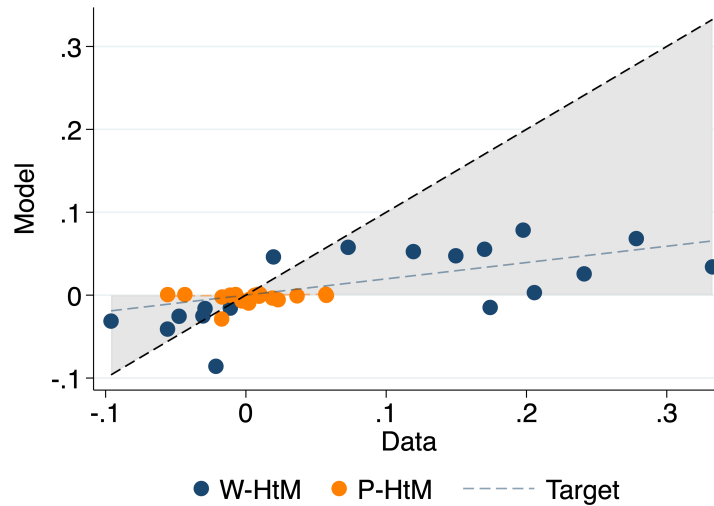
Figure 12: Switching One Component at a Time



Notes: Each dot represents the difference in the share of HtM households (wealthy in blue, poor in orange) between a given country and Spain (baseline). In the horizontal axis is the difference in the data. In the vertical axis is the difference obtained in the model by changing a subset of parameters of the income process to those estimated for the corresponding country (leaving all others unchanged). In panel (a) we change the parameters of the transitory component of stochastic income (z_1), in (b) those of the persistent component (z_2), in (c) those for the pension scheme, and in (d) those of the age profile of income ($\mu_z(t)$). The shaded area between the horizontal axis and the 45-degree line represents the case where the change gets the model closer to reproducing the difference in the data between Spain and the corresponding country.

The differences in parameters of the transitory shock also affect the shares of W-HtM and P-HtM significantly. They seem to have a small positive contribution in increasing the W-HtM share for countries with higher shares, but it also pushes up the share for the countries that have a lower share of W-HtM than Spain. The parameters of the persistent shock and the pension scheme do not seem to contribute much to explain the differences. They are related to low-frequency events that therefore appear to have little influence on

Figure 13: Financial Conditions



Notes: Each dot in the figure represents the difference in the share of HtM households (wealthy in blue and poor in orange) between a given country and Spain (baseline). In the horizontal axis is the difference observed in the data. In the vertical axis is the difference we obtain in the model when we change interest rate spreads and the borrowing limit to the corresponding country. The shaded area between the horizontal axis and the 45-degree line represents the area where these changes moves the model closer to reproducing the difference observed in the data between Spain and the corresponding country.

the demand for liquidity.

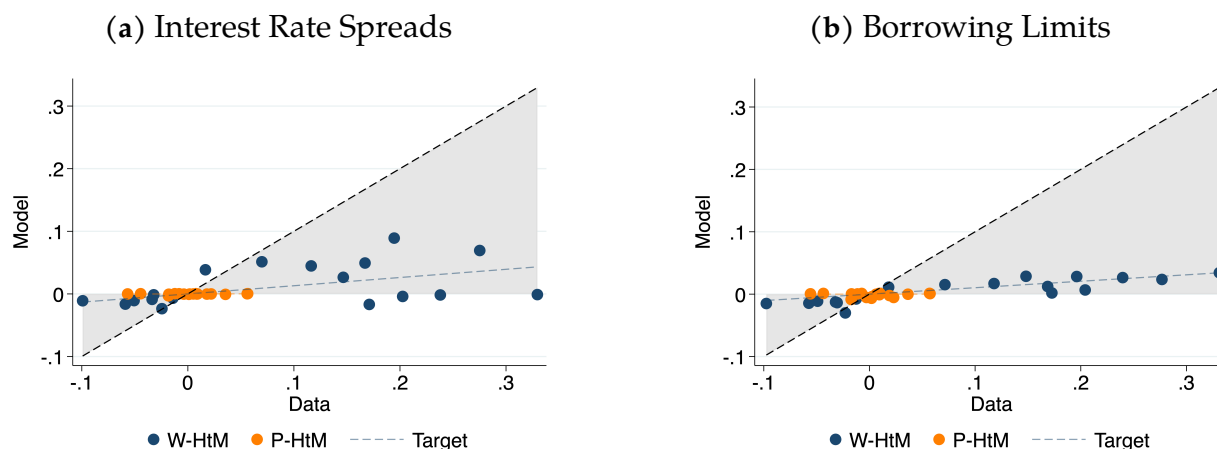
5.2 Financial Conditions

Falling into HtM status is naturally related to the household's demand for liquidity and the ability to borrow when hit by negative shocks. The interest rate (for saving and borrowing) and borrowing limits are at the core of the liquidity demand decision. We focus on these two financial frictions as they are more easily relatable with data counterparts. To measure interest rate spreads we used data from Eurostat.¹⁶ In the case of borrowing limits we used as proxy 30% of average income, which is a usual measure taken by the literature.

Figure 13 shows that differences in the financial conditions considered contribute to explaining the differences in the share of W-HtM. Considering all countries in the sample,

¹⁶We take an average from 2003 to 2019 of deposit and loans for consumption rates, both with lower than a year maturity

Figure 14: Individual Financial Conditions



Notes: Each dot represents the difference in the share of HtM households (wealthy in blue, poor in orange) between a given country and Spain (baseline). In the horizontal axis is the difference in the data. In the vertical axis is the difference obtained in the model by changing a one of the financial conditions. In panel (a) we change the interest rate spreads and in (b) the borrowing limit. The shaded area between the horizontal axis and the 45-degree line represents the case where the change gets the model closer to reproducing the difference in the data between Spain and the corresponding country.

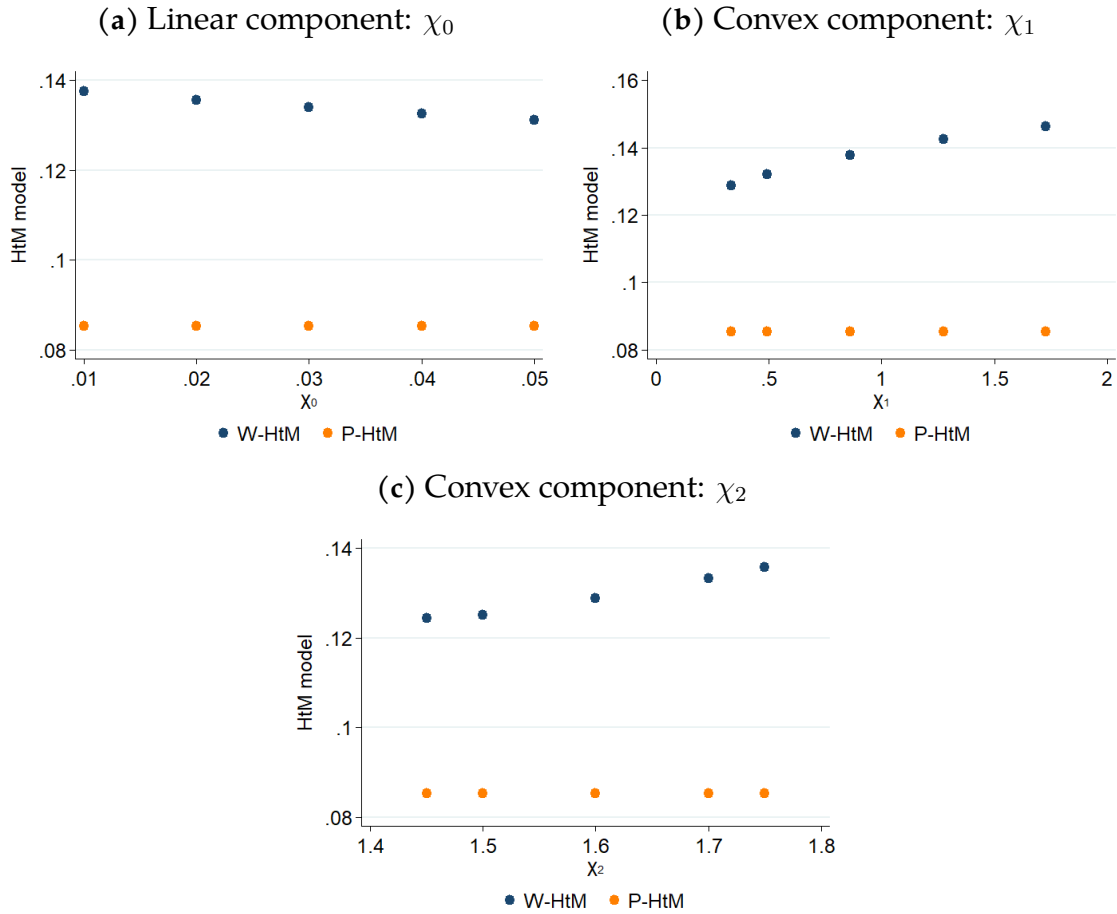
financial conditions by themselves are able to explain 68% of the W-HtM differences on average. Their explanatory power is stronger among countries with lower W-HtM shares, and gradually declines as the share increases. Interestingly P-HtM shares are barely affected, with simulations only explaining 17% of the differences.¹⁷

When we look at each of the frictions individually we find that both contribute to explaining the differences in the share of HtM. Interest rate spreads explain 31% of the difference and the borrowing limit a 15%. With respect to W-HtM spreads are able to reproduce 34% and the borrowing limit 29%. The results indicate that there is a complementarity between the two frictions, as the explanatory power when both are included is higher than the sum of the individual exercises.

Illiquid transaction costs. These costs are meant to capture, in reduced form, frictions in the housing and mortgage markets. Moreover, there is evidence of frictions in these mar-

¹⁷The percentages reported are computed as the average of the ratio between the distance of the simulated share in the model to the baseline calibration and the distance in the data between each given country and Spain.

Figure 15: HtM sensitivity to Adjustment Costs



Note: This figure plots the share of wealthy (blue dots) and poor (orange dots) hand-to-mouth agents that correspond to different values of χ_0 (in panel a), χ_1 (in panel b), and χ_2 (in panel c), leaving all other parameters unchanged.

kets being correlated with HtM shares.¹⁸ However, we do not have a clear data counterpart of the model parameters. Alternatively, we simulate the model modifying the value of the calibrated baseline parameters, keeping the rest of the model fixed. Figure 15 shows the results.

Adjusting these parameters mainly impacts the W-HtM share, with minimal effect on the P-HtM. A higher value of χ_0 , the parameter governing the linear component of the adjustment cost, expands the region where households refrain from transferring resources to or

¹⁸see Figure 5 in the appendix

from the illiquid account. This has a mixed impact on HtM shares. In some scenarios, households opt against depositing liquid balances into the illiquid account due to costs, keeping positive liquid balances and thus reducing the HtM share. Conversely, in other cases, households abstain from withdrawing from the illiquid account, keeping zero balances and thus increasing the HtM share. The first force seems to dominate in the model, as panel (a) shows that the share of W-HtM is slightly decreasing with χ_0 .

The convex part of the adjustment cost function is governed by two parameters, χ_1 and χ_2 . When these parameters are larger, households tend to make larger but less frequent deposits and withdrawals. Again, this can have ambiguous effects on the share of HtM households, as some may hold more liquid balances compared to lower χ_1 or χ_2 , while others may hold less. The second effect appears to be stronger, as panels (b) and (c) show that larger parameter values are associated in the model with a higher share of HtM households.

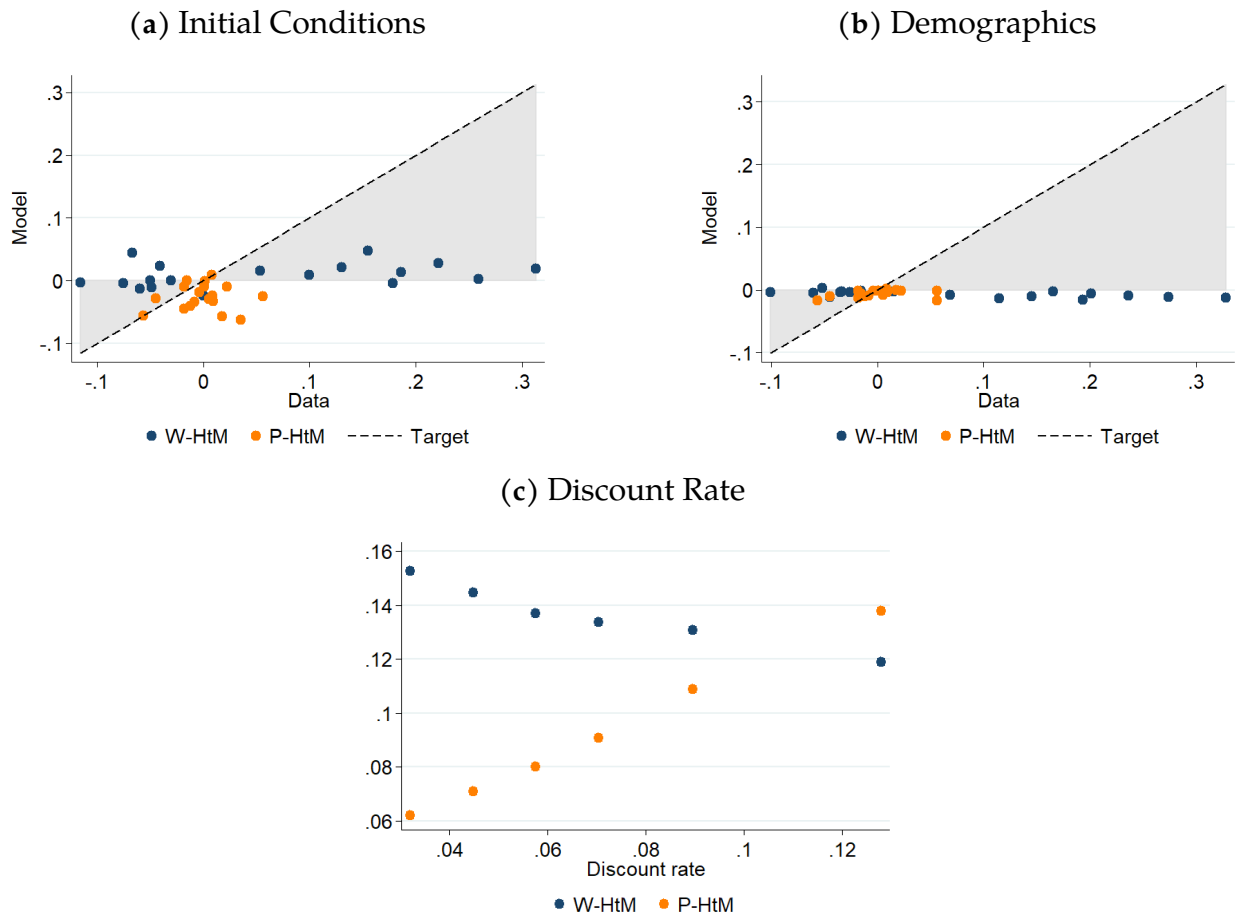
Overall, liquid spreads and borrowing limits have a key role in explaining differences in the HtM shares among countries with lower than 35% of HtM, by explaining differences in the share of W-HtM. They also contribute to raise the shares of HtM among countries with larger shares, but their contribution falls. Moreover, there is scope for frictions in the housing and mortgage markets to contribute to differences in the shares of HtM across countries. However the lack of a clear data counterpart for our model does not allow us to quantitatively assess its explanatory power. Explicitly modeling the housing market is a promising avenue for future research.

5.3 Initial Conditions, Demography, and Discount Rate

We turn now to analyze through the lens of the calibrated model whether other drivers have the potential to explain some of the differences in HtM shares. Households enter the model with a certain amount of assets that are calibrated to match the initial joint distribution of income and assets in the data. Panel (a) in Figure 16 shows the results of modifying the initial asset distribution to the one of each country. Initial conditions are relevant espe-

cially for countries with shares of W-HtM higher than Spain, as they increase the simulated share for all but one country. For most of the countries with a lower share of W-HtM, the initial conditions play a minor role as they are likely to be similar to the Spanish one.

Figure 16: HtM Shares and Other Potential Drivers



Note: In panels (a) and (b) each dot represents the difference in the share of HtM households (wealthy in blue, poor in orange) between a given country and Spain (baseline). In the horizontal axis is the difference observed in the data. In the vertical axis is the difference we obtain in the model when we change the initial asset distribution (panel a) and the demographic structure (panel b) to those for each country (leaving all other parameters unchanged). The shaded area between the horizontal axis and the 45-degree line represents the case where the change gets the model closer to reproducing the difference in the data between Spain and the corresponding country. Panel (c) plots the share of wealthy (blue dots) and poor (orange dots) hand-to-mouth agents that correspond to different values of the discount rate leaving all other parameters unchanged.

Life expectancy and survival probabilities also differ between countries and can shape the households' incentives to accumulate wealth. However, we find that modifying them has

a limited effect on the share of HtM, as can be seen in panel (b). Finally, we explore in panel (c) how the share of HtM changes when we modify the discount rate. When agents are more impatient, that is when the discount rate is higher, the share of W-HtM falls, and the share of P-HtM increases. For the W-HtM there are opposing forces: there are fewer households at the zero kink but more at the credit limit. On the other hand, for the P-HtM the share of households at both kinks increases. The strong response of the P-HtM share is at odds with the fact that it does not differ so much in the data, making differences in preferences across countries not likely to be important drivers of the difference in HtM shares.

6 Conclusions

A burgeoning literature is studying the effects of incorporating heterogeneity on the efficacy and transmission channels of fiscal and monetary policies. This literature highlights the role of HtM households, which have relatively large MPCs, as one of the key distributional moments that shape the aggregate response of the economy to these policies. Countries differ in the share of households that are HtM, but little is known about what drives these cross-country differences. This paper begins to study the drivers of these cross-country differences by analyzing the role of income and other potential drivers.

We first document significant heterogeneity in the share of HtM across countries in Europe using the Eurosystem Household Finance and Consumption Survey (HFCS). We find that differences in the share of wealthy HtM (households who only hold illiquid assets like real estate) are behind the HtM differences, as poor HtM represent a similar proportion across countries.

We build a life-cycle model with incomplete markets and idiosyncratic risk in the tradition of ?, ?, and ?. Households have access to two assets, a liquid and an illiquid one, and face three financial frictions: a transaction cost to access the illiquid account, a spread between the saving and borrowing rate in the liquid assets, and a borrowing limit. As a consequence

of this structure, the share of HtM is determined endogenously in the model.

We calibrate the model for Spain, a country that ranks in the middle of the HtM country-level distribution. With the model, we study the role of differences in income as drivers of the differences in HtM shares. Decomposing income between a deterministic age profile, a persistent shock, a transitory shock, and a pension scheme, we analyze the role of the level and income risk over the life-cycle of households in determining the share of HtM. We find that the share of W-HtM is most responsive to changes in the transitory component of risk, while the share of P-HtM is most responsive to the age profile. However, when we switch all parameters of the income process at the same time, we find that they fall short of matching countries' fraction of HtM agents.

Finally, we further investigate from the lens of our calibrated model other potential drivers and find that differences in financial frictions (borrowing limits and interest rates) can play a major role and are a promising line for future research on this question. We have focused on quantifying the contribution of each potential driver in isolation, but the existence of complementarities among them should also be taken into account.

Our findings have relevant policy implications, in particular with respect to the regulations and policies related to households borrowing costs. Understanding why borrowing costs are higher in countries with higher shares of HtM, and taking actions to reduce those can have important macroeconomic stabilization effects. Through reducing the share of HtM, and the average marginal propensity to consume of the economy, countries can become more resilient to shocks and foster economic growth. As a consequence pro-competitive policies in the banking sector like open banking and the defense of the consumer of financial services, have important macroeconomic stabilization effects that are not usually taken into account.

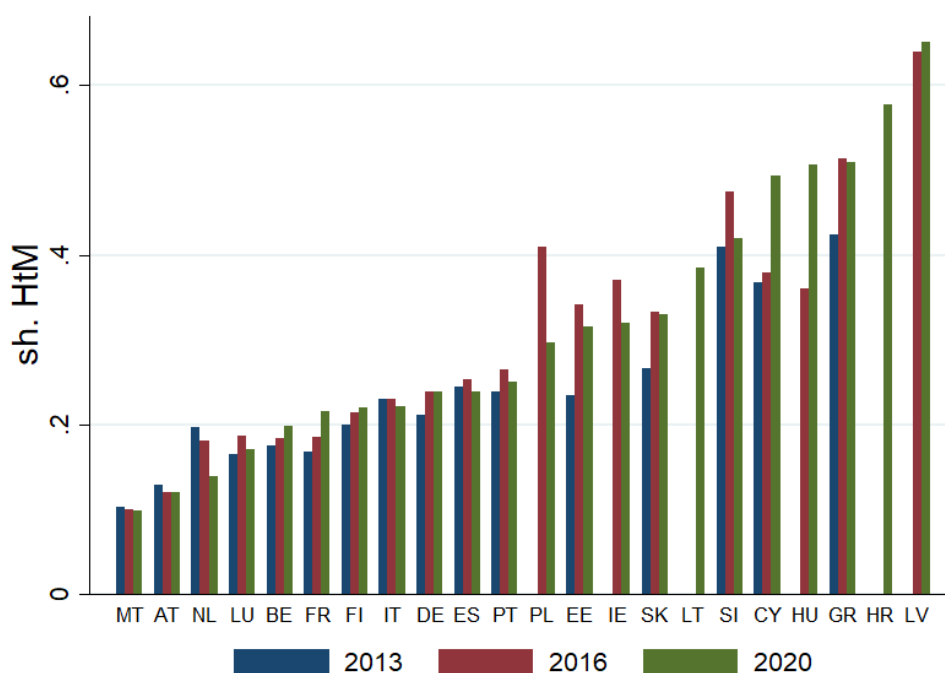
Appendix

A Measurement of HtM Households

A.1 The Share of HtM in Europe over Time

Figure 17 show the share of HtM households over time for all countries in the sample. We exclude the last wave, released in 2023, because it was conducted during the pandemic.

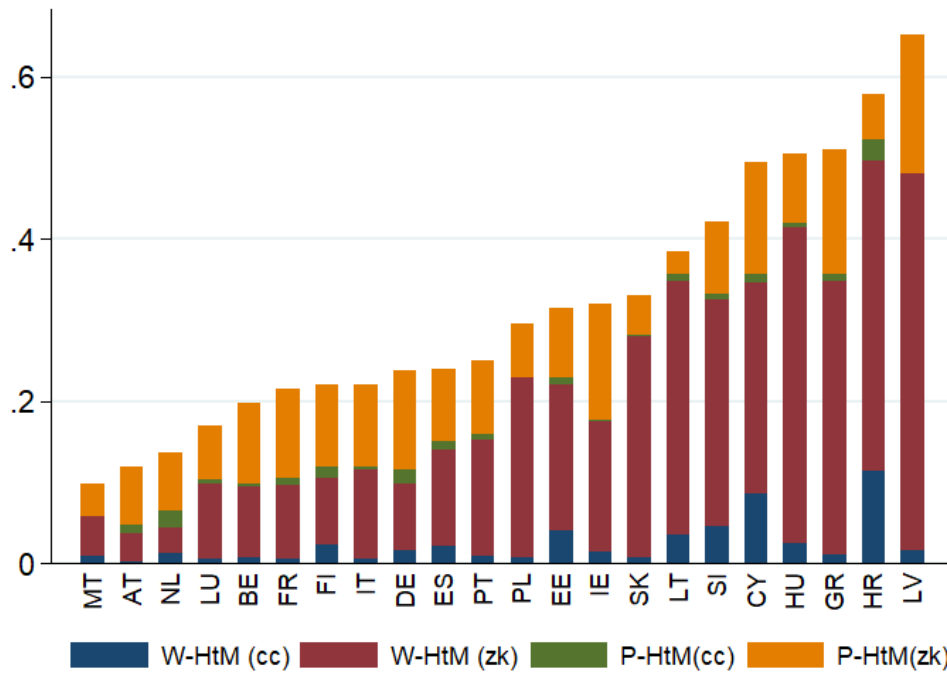
Figure 17: HtM by country over time



A.2 Decomposition of HtM Households

Households are classified as HtM if their net liquid wealth is close to zero (related to the zero kink in the Euler equation) or if it is too negative (implying that they are close or have hit the borrowing limit). Figure 18 shows the shares of wealthy and poor HtM in the zero kink (zk) or who are credit constrained (cc), according to our baseline measure. In all countries, the majority of HtM households (both poor and wealthy) are at the zero kink.

Figure 18: Decomposition by kinks



A.3 Robustness on the Measurement of HtM Households

In this section we test how robust the measurement of HtM is to modifications in the assumptions we make. First, we look at changes in the classification of liquid and illiquid assets. We change the classification of retirement accounts from illiquid to liquid for households older than 60, we include cars and other miscellaneous items (like jewelry and art pieces) as illiquid assets and we change the classification of directly held stocks and mutual funds from liquid to illiquid. Second, we change the assumed frequency of income payments from monthly to both weekly and bi-weekly. Third, we add consumption loans as liquid debt and we increase the credit limit, which does not have much impact since most households are at the zero kink. Finally, we add additional cash balances to all households equal to 15% of their income. This is to account for evidence that cash may be used more intensely in poorer countries, which is not captured in our imputation of cash.

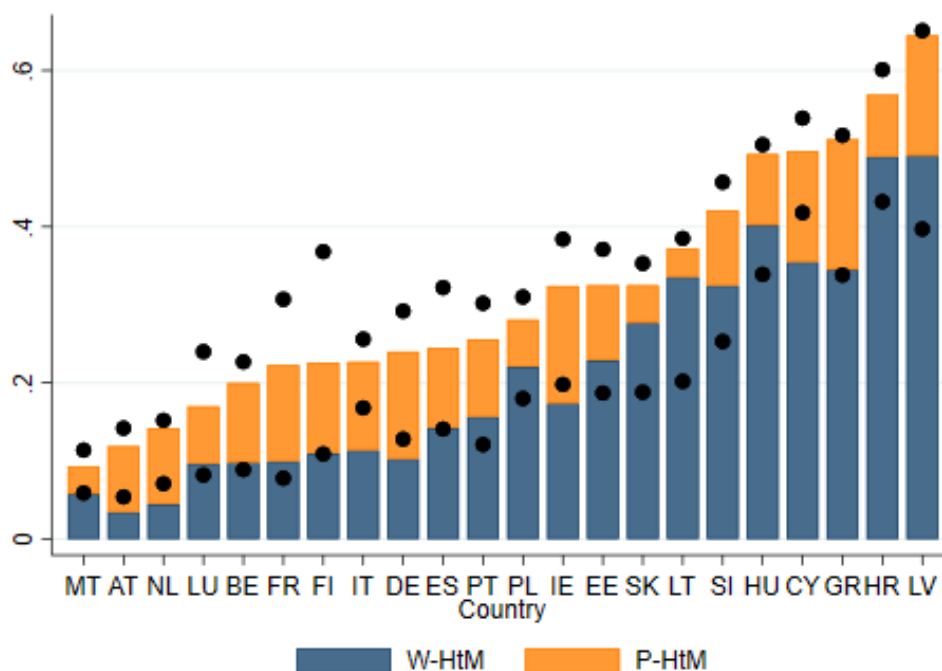
Table 8 shows the results for all countries and each robustness exercise. Figure 19 compares the share of HtM under the baseline assumptions and the minimum and maximum coming from the robustness exercise. Overall, the differences in HtM shares across countries are maintained.

Table 8: Robustness checks

	IT	ES	MT	AT	BE	CY	EE	IE	FI	FR	GR
Baseline	.227	.245	.093	.119	.2	.497	.325	.324	.226	.223	.512
1 year income credit limit	.215	.214	.086	.105	.187	.418	.276	.306	.191	.208	.495
Businesses as illiquid assets	.23	.243	.098	.12	.2	.485	.326	.326	.225	.221	.506
Consumption loans	.256	.322	.112	.142	.227	.519	.371	.384	.368	.307	.517
Direct as illiquid assets	.237	.249	.114	.122	.211	.509	.327	.348	.253	.23	.514
Excludes cc puzzle hh's	.229	.243	.095	.12	.208	.539	.347	.356	.228	.222	.517
Bi-weekly pay period	.186	.177	.073	.073	.134	.448	.243	.25	.157	.132	.41
Other valuables as illiq. assets	.227	.245	.092	.12	.201	.497	.325	.323	.226	.222	.512
Ret. acc. as liquid for 60+	.227	.245	.093	.119	.2	.497	.325	.324	.226	.223	.512
Vehicles as illiquid assets	.227	.245	.093	.119	.2	.496	.326	.324	.226	.222	.511
Weekly pay period	.168	.141	.059	.054	.089	.419	.187	.198	.109	.078	.338
Additional Cash	.207	.207	.081	.1	.172	.477	.313	.298	.199	.186	.451

	HR	HU	LT	LU	LV	NL	PL	PT	SI	SK	DE
Baseline	.569	.493	.372	.17	.645	.142	.281	.256	.421	.325	.24
1 year income credit limit	.447	.459	.329	.161	.626	.111	.273	.24	.366	.315	.206
Businesses as illiquid assets	.575	.485	.385	.172	.641	.143	.282	.258	.422	.332	.241
Consumption loans	.593	.497	.384	.24	.626	.147	.31	.302	.457	.353	.292
Direct as illiquid assets	.588	.501	.38	.176	.646	.152	.283	.259	.432	.33	.247
Excludes cc puzzle households	.601	.505	.378	.165	.651	.134	.31	.26	.431	.335	.245
Bi-weekly pay period	.482	.404	.271	.116	.526	.093	.213	.176	.319	.245	.172
Other valuables as illiq. assets	.569	.493	.372	.17	.645	.142	.281	.256	.422	.325	.24
Ret. acc. as liquid for 60+	.569	.493	.372	.17	.645	.142	.281	.256	.421	.325	.24
Vehicles as illiquid assets	.57	.493	.372	.17	.645	.142	.281	.256	.421	.325	.24
Weekly pay period	.432	.339	.202	.082	.397	.071	.18	.121	.253	.188	.128
Additional Cash	.539	.452	.331	.148	.607	.121	.248	.214	.379	.288	.213

Figure 19: HtM sensitivity



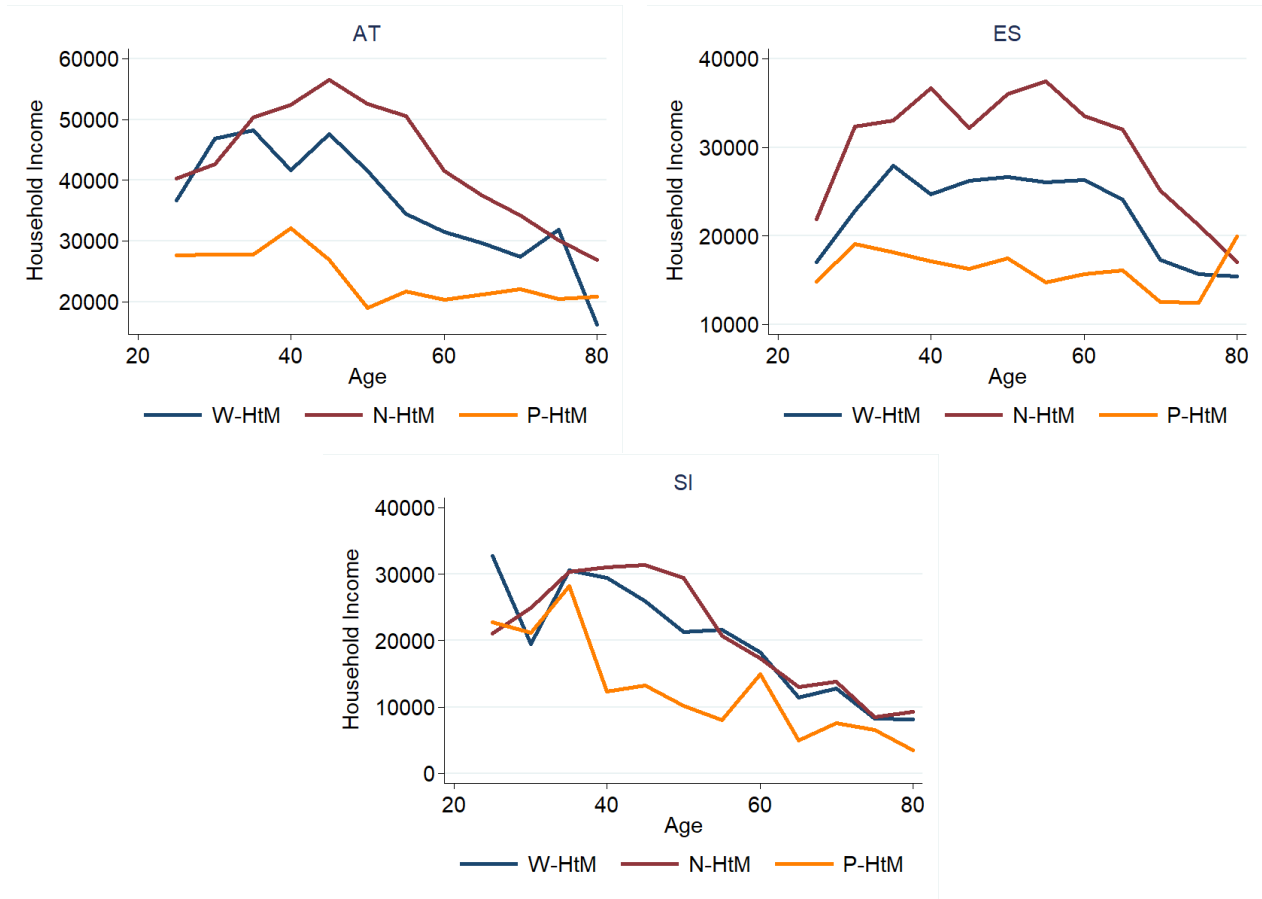
B Descriptive Statistics

Income over the Life Cycle Figure 20 compares, for each of the three countries, the life-cycle profile of income of the non-HtM, the P-HtM and the W-HtM. For most countries, households who fall into W-HtM status have on average and income that is in between that of P-HtM and N-HtM. For other countries, like Slovenia (but also true for low HtM countries like the Netherlands), the average income of the W-HtM is similar to that of the N-HtM.

Balance Sheets Figure 21 shows the balance sheet structure of W-HtM and Non-HtM divided by assets and liabilities. Households hold most of their wealth in real estate. This is particularly true for W-HtM households, whereas Non-HtM households own a more diversified portfolio, which includes by definition liquid assets like cash, deposits and savings accounts, but also a larger share of other illiquid assets like cars. Notably, the share of wealth held in retirement accounts is negligible for all countries, meaning that public pension plans are the most important source of income during retirement. Countries in the figure are ordered according to the share of W-HtM they have. Interestingly, Non-HtM households in countries with larger shares of W-HtM, hold a lower share of liquid assets in their portfolio which could imply a higher risk of becoming HtM.

Comparing the debt structure, panels (c) and (d) show that the structure of debt is very

Figure 20: Income by HtM



similar for W-HtM and N-HtM households, with mortgages being the main instrument used by both types of households¹⁹.

Tables 9 and 10 show data on the balance sheet composition of W-HtM and Non-HtM for three sample countries (Austria, Spain, and Slovenia).

¹⁹This does not imply that the levels of debt are similar, only the composition across instruments is. See appendix B where we show extended comparisons analyzing the balance sheets and other characteristics of the households like marital status, educational achievement and sources of income across HtM types. Additionally, appendix B.1 presents comparisons across countries offering suggestive evidence on some of the potential determinants behind the difference in the share of HtM.

Figure 21: Balance sheet structure



Table 9: Balance sheet of W-HtM

	AT	ES	SI
Monthly Income	3432	2018	1876
Net Liq. Wealth	516	-733	-6
Cash and deposits	545	612	275
Shares and bonds	2	7	11
Saving Accounts	440	20	39
Broad Credit Card	470	1373	331
Net Illiq. Wealth	190005	144099	125240
Houses	206804	167777	122157
Other properties	0	1844	629
Cars (net)	8164	4841	5968
Retirement accounts	0	0	0
Life Insurance	1213	5058	1404
Residential debt	26175	34598	4604
Home equity LOC	440	823	314

Table 10: Balance sheet of N-HtM

	AT	ES	SI
Monthly Income	3914	2446	1731
Net Liq. Wealth	38030	40531	13148
Cash and deposits	4185	18451	5867
Shares and bonds	6334	15228	3038
Saving Accounts	27631	6945	4354
Broad Credit Card	119	92	111
Net Illiq. Wealth	178173	242118	136759
Houses	178447	244019	131149
Other properties	3012	10396	1993
Cars (net)	10986	7390	7197
Retirement accounts	0	0	0
Life Insurance	3108	8138	1900
Residential debt	17353	27721	5464
Home equity LOC	27	104	15

Other demographics Figure 22 compares the share of W-HtM among households whose head is married with the share among the single/never married at different ages . Although there are differences between these two groups, there is no systematic correlation between marital status and being HtM. We also look at differences in HtM status by education level in Figure 23. The share W-HtM households is clearly decreasing with education level in all countries (once we ignore those whose higher level of education is Primary school which is noisy).

Figure 22: W-HtM by marital status

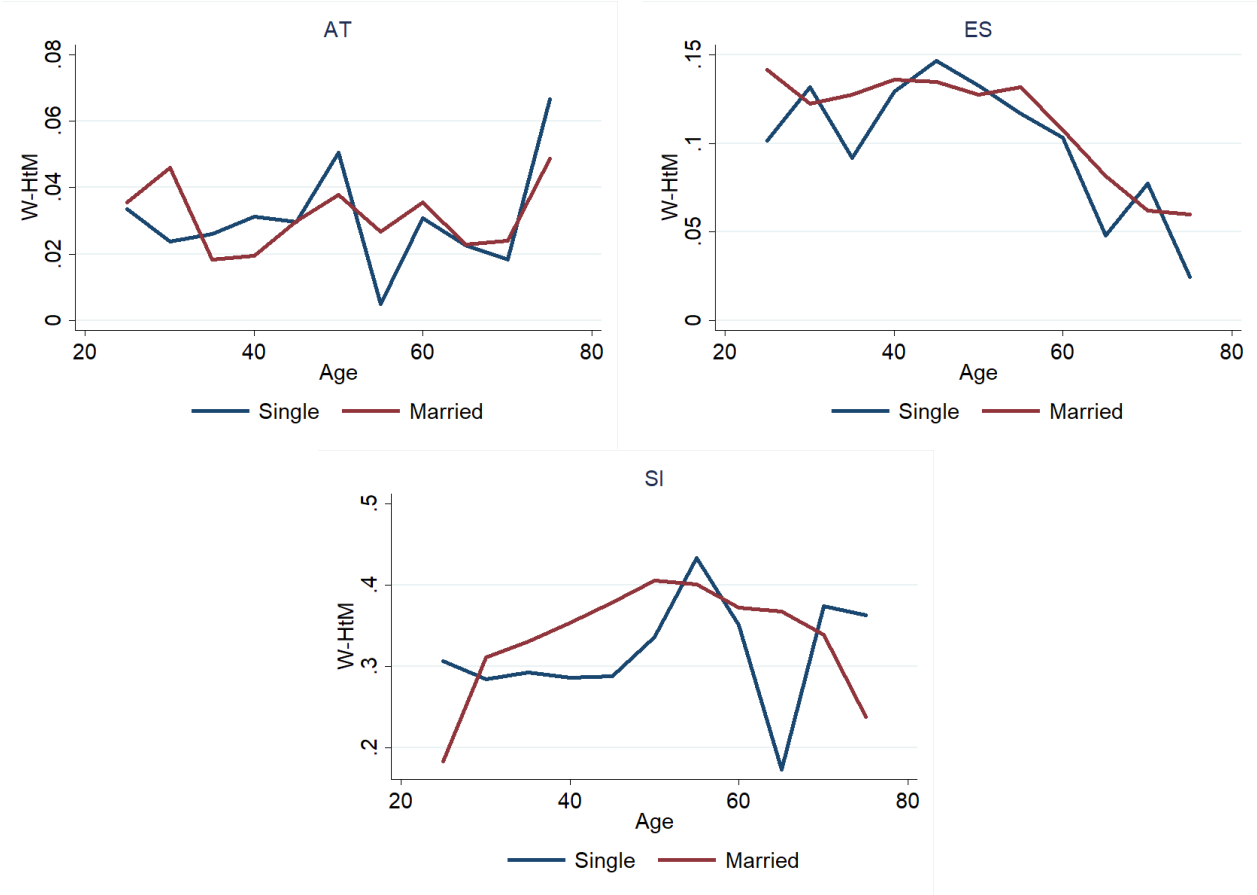


Figure 23: HtM by education



Table 11: Correlations

	HtM	WHtM	PHtM
(log) GDP	-0.659**	-0.737***	0.103
Financial Development	-0.609**	-0.698***	0.172
Expected pension	-0.608*	-0.616*	-0.256
Var income (45)	0.513*	0.474*	0.317
Var 1 yr changes	-0.0305	-0.0593	0.115
Income changes less than 10%	-0.290	-0.336	0.117
Income growth 25 - 55	0.303	0.284	0.166
Observations	22	22	22

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 12: Regression

	HtM	WHtM	PHtM
(log) GDP	-0.155*** (0.0438)	-0.145*** (0.0420)	-0.00963 (0.0124)
Financial Development	-0.0552 (0.0956)	-0.0741 (0.0864)	0.0376 (0.0452)
Observations	72	72	72

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

B.1 Cross-Country Comparisons

Table 11 reports the cross-country correlations of the share of HtM, W-HtM, and P-HtM for the different variables explored in section 2.5. Table 12 reports the results of regressing countries' share of HtM on its log GDP and the IMF's Financial Development Index using all waves. Only the coefficient for GDP is significant (and negative for the total HtM and the wealthy HtM), while the coefficient for financial development is negative but not statistically significant.

Retirement. To provide more evidence on the link between retirement and HtM status we regress W-HtM status (a dummy that equals one when households are W-HtM) on expected replacement rates, household fixed effects and other controls. Table 13 shows the results. In the first column we include all households, in the second only those that hold positive illiquid balances, in the third only young households (with the head being younger than 35 years old) and in the fourth column only older households (older than 50 years old). Higher expected pensions are associated with a lower probability of becoming W-HtM. This effect is stronger for households whose head is close to retirement and therefore have more information about their future retirement income.

Figure 24: Variance of income at different ages

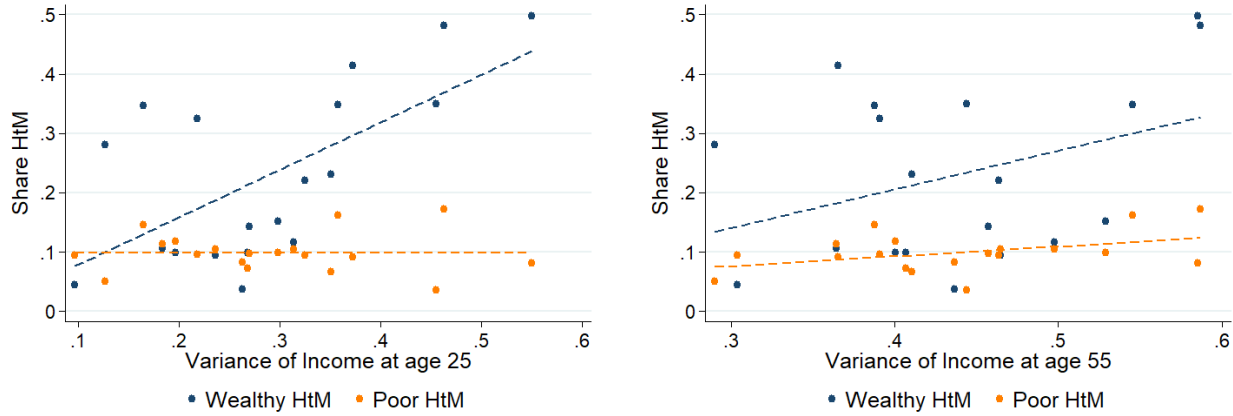


Figure 25: Variance of income changes

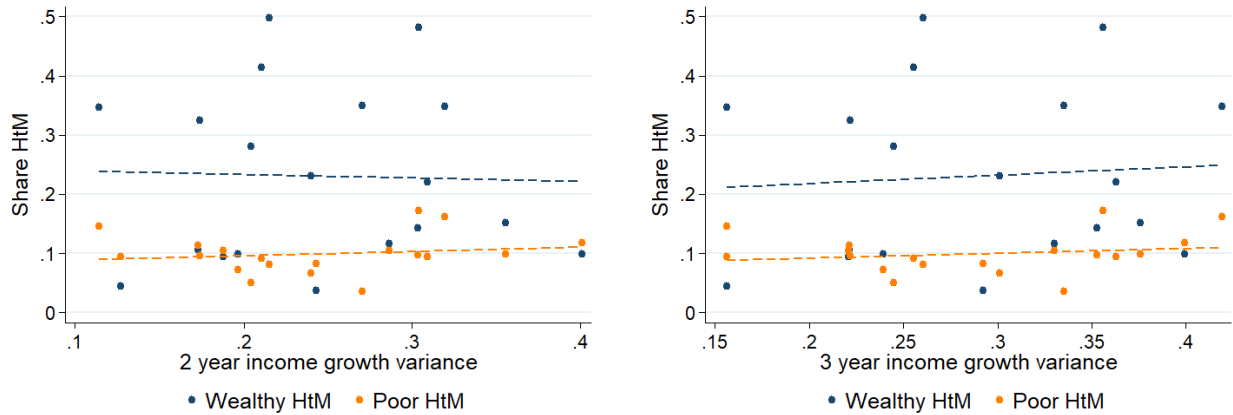


Table 13: W-HtM status

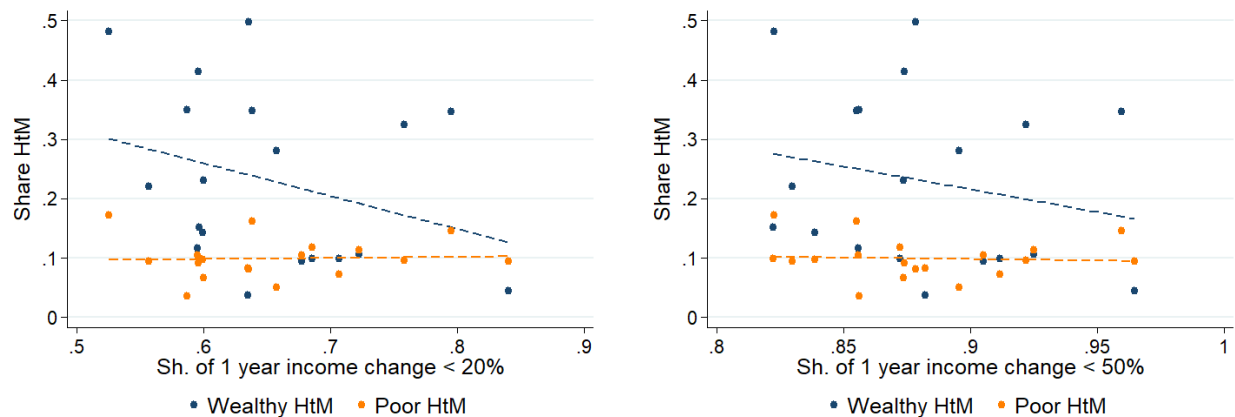
	(1) All	(2) Wealthy	(3) Young	(4) Old
(Log) labor income	0.00100 (0.0114)	-0.00986 (0.0136)	-0.0170 (0.0241)	0.00855 (0.0210)
Mortgage payments over income	0.111 (0.0700)	0.109 (0.0746)	-0.0566 (0.229)	0.185* (0.107)
Owns house	0.0834*** (0.0267)	0.0242 (0.0349)	0.156** (0.0682)	0.0637 (0.0396)
Expected replacement rate	-0.0356 (0.0259)	-0.0511* (0.0304)	0.179** (0.0760)	-0.0929** (0.0391)
Observations	53217	45447	5489	25493

Standard errors in parentheses

Age, Country FE included

* p<0.10, ** p<0.05, *** p<0.01

Figure 26: Share of 1 year income changes close to zero



C Estimation of the Income Process

Figure 27 shows for all countries how the estimation of the income process matches the data in terms of the variance of income at different ages. The fit for the other moments targeted in the estimation are summarized in Figure 28. Most of the moments match the data closely with the exception of the first order autocorrelation of income at different horizons, which is lower than in the data in most cases.

Figure 27: Model fit

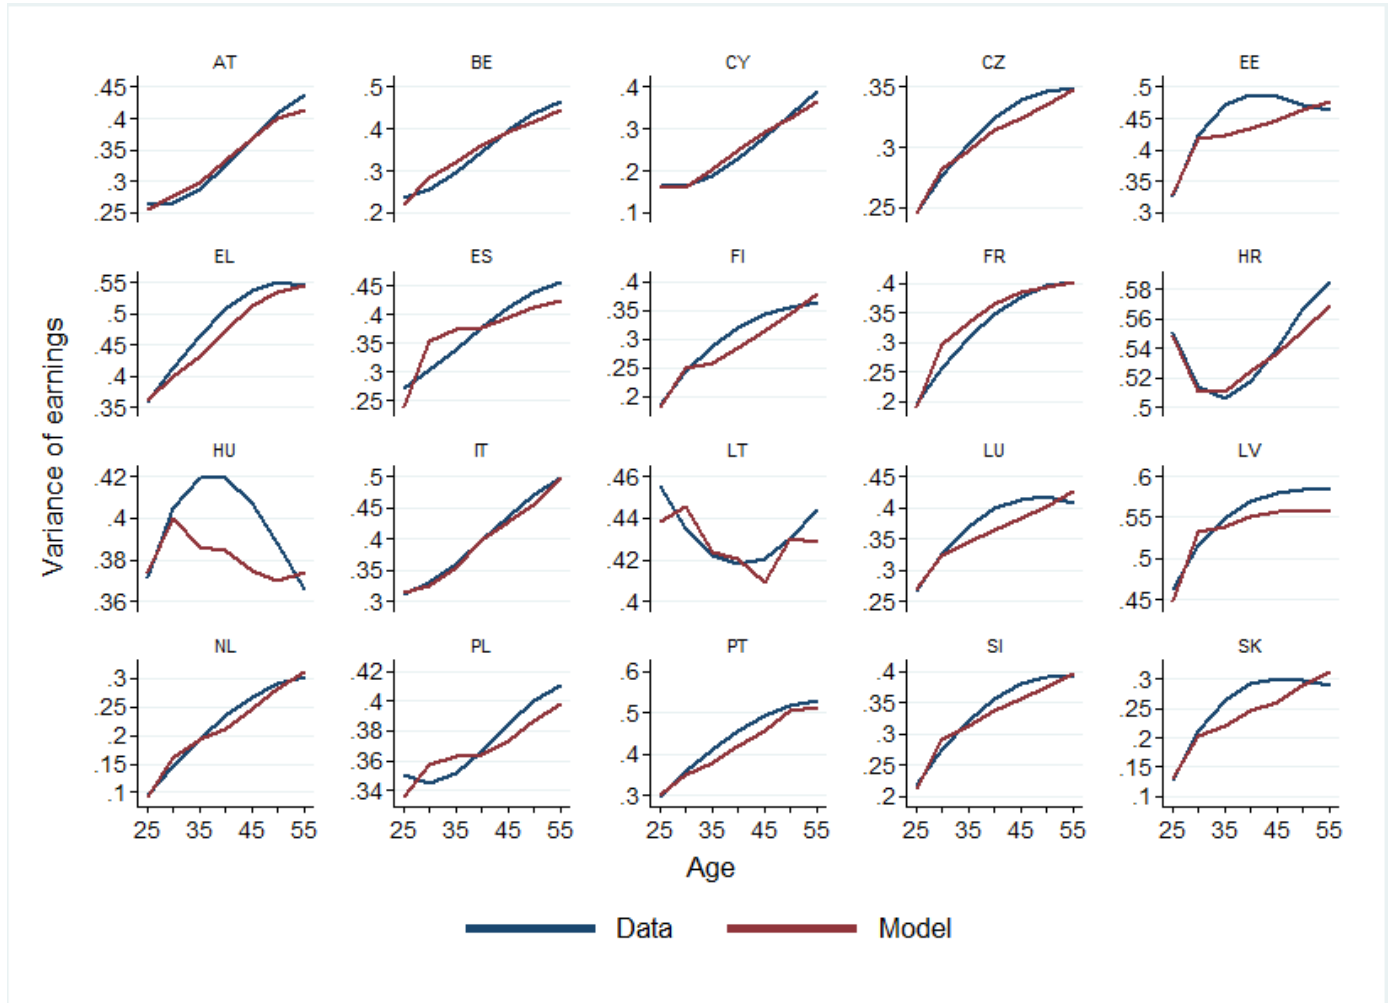
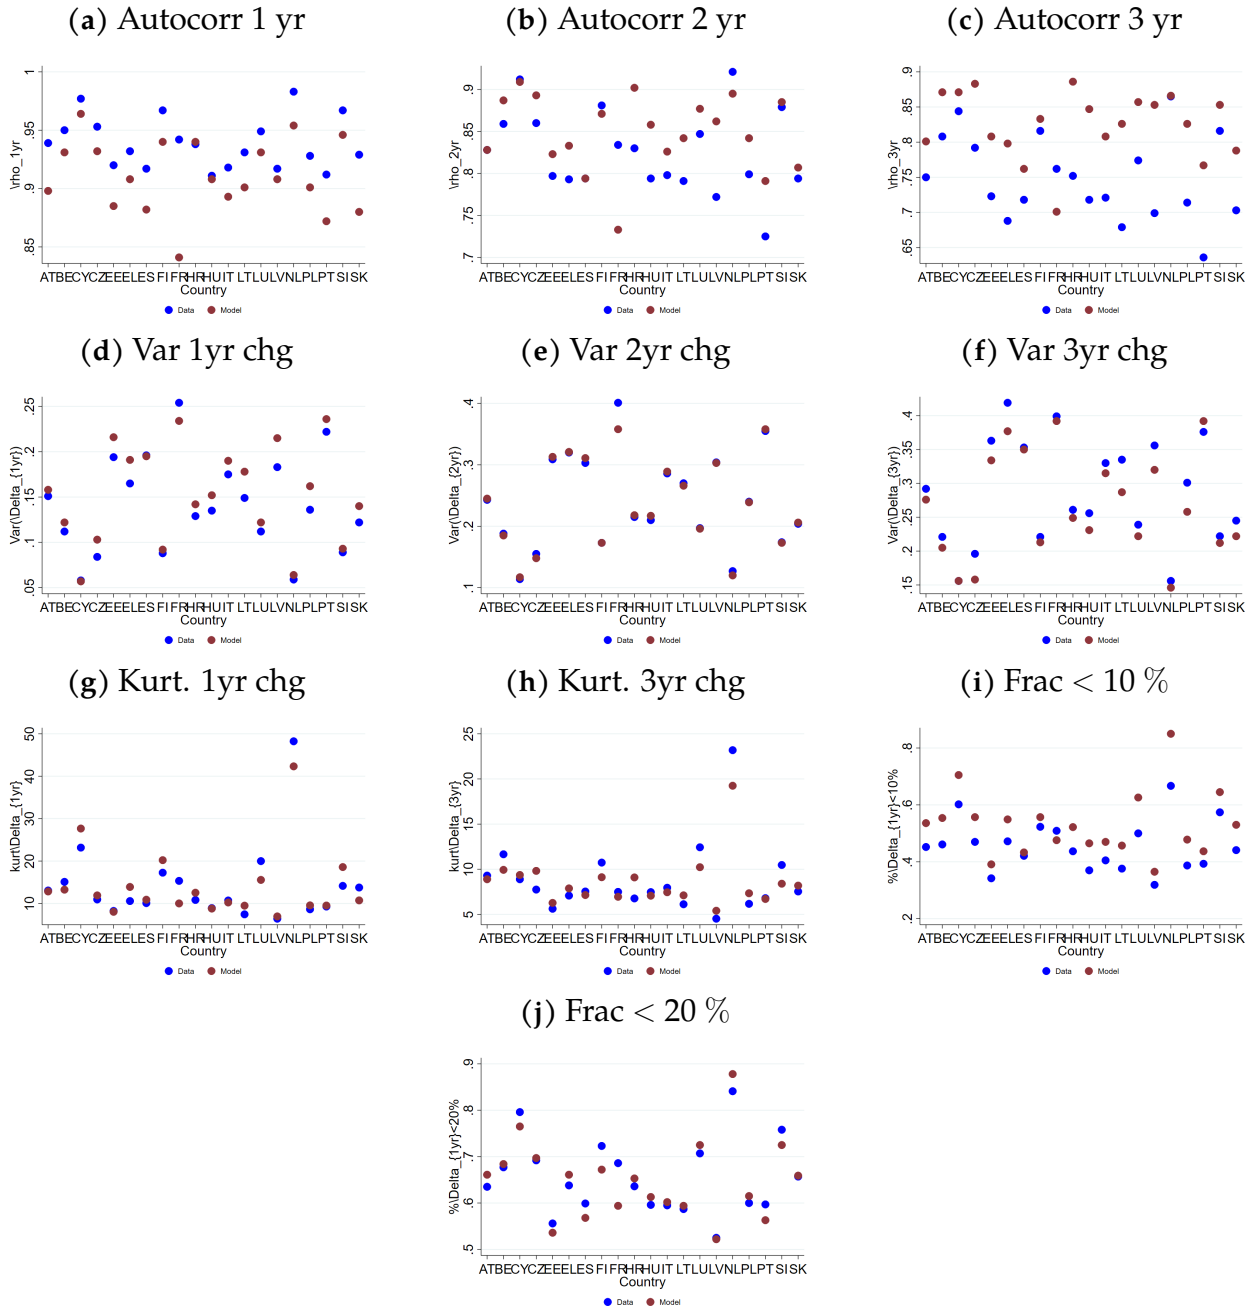


Figure 28: Model's fit



The estimated parameters for all countries are presented in table 14.

Table 14: Estimated Parameters Income Process

Shock	Parameter	AT	BE	CY	CZ	EE	EL	ES	FI	FR	HR
Transitory	λ_1	0.355	0.357	0.126	0.365	0.591	0.294	0.324	0.304	0.404	0.388
Transitory	σ_1	0.862	0.775	0.707	0.736	0.837	0.963	1.018	0.627	0.997	0.784
Transitory	β_1	0.897	0.882	0.280	0.883	0.857	0.850	0.884	0.311	0.991	0.901
Persistent	λ_2	0.008	0.006	0.007	0.002	0.004	0.017	0.002	0.004	0.008	0.011
Persistent	σ_2	1.635	1.500	1.278	1.586	1.662	1.466	1.371	1.812	1.577	1.384
Persistent	β_2	0.045	0.010	0.004	0.004	0.014	0.054	0.0004	0.027	0.062	0.018

Shock	Parameter	HU	IT	LT	LU	LV	NL	PL	PT	SI	SK
Transitory	λ_1	0.485	0.421	0.468	0.251	0.647	0.062	0.449	0.457	0.210	0.379
Transitory	σ_1	0.781	0.881	0.829	0.876	0.817	1.157	0.812	0.946	0.763	0.830
Transitory	β_1	0.884	0.851	0.846	0.847	0.882	0.805	0.846	0.907	0.451	0.929
Persistent	λ_2	0.004	0.004	0.005	0.006	0.008	0.004	0.004	0.008	0.008	0.004
Persistent	σ_2	1.392	1.791	1.476	1.593	1.032	1.633	1.639	1.564	1.340	1.477
Persistent	β_2	0.012	0.012	0.018	0.019	0.004	0.012	0.015	0.024	0.015	0.010

D Additional Results

Figure 29: HtM and income process parameters

