

How Do Households Respond to Income Shocks?*

Dirk Krueger Egor Malkov Fabrizio Perri
University of Pennsylvania Bates White Federal Reserve Bank
CEPR and NBER of Minneapolis and CEPR

August 2024

Abstract

We use panel data from the Italian Survey of Household Income and Wealth from 1991 to 2016 to document what components of the household budget constraint change in response to shocks to household labor income, both over shorter and over longer horizons. Consumption and wealth responses are informative about the household consumption (or savings) function and thus about what class of consumption-savings model best describes the data. Empirically, we first show that shocks to labor income are associated with negligible changes in transfers and non-labor income components, modest changes in consumption expenditures, and large changes in wealth. To understand the wealth response we then split households into a sample that does not own business or real estate wealth, and a sample that does. For the first group, we find that consumption responses are more substantial (and increasing with the horizon of the income shock) and wealth responses are much smaller (and mildly increasing with the income shock horizon). Turning to theory, we argue that for this group, a simple extension of the standard permanent income hypothesis (PIH) consumption function that allows for partial insurance against even permanent income shocks explains the consumption and wealth responses well, both at short and long horizons. For the second group with business wealth or real estate wealth the standard framework cannot explain the large changes in wealth associated with income shocks. We conclude that models which include shocks to the value of household wealth are necessary to fully evaluate the sources and consequences of household resource risk.

JEL CODES: D91, E21

KEYWORDS: Consumption, Liquid Wealth, Income Shocks, Partial Insurance, PIH

*We thank Ctirad Slavik and Jacek Rothert for excellent research assistance, participants at many seminars and conferences and our discussant Martin Holm for many helpful suggestions and his written discussion (Holm 2023), and the NSF for financial support under grant SES-0820494. This is a significantly revised and updated version of Krueger and Perri (2011) with the same title. The opinions expressed represent those of the authors, and do not represent the views or opinions of the Federal Reserve Bank of Minneapolis, the Federal Reserve System, Bates White, LLC or of other Bates White employees or affiliates.

1 Introduction

What do households do when confronted with a shock to their labor incomes? To answer this question empirically, we use a panel data set from Italy that contains detailed information about household income, consumption, and wealth, in order to document how the different components of the household budget constraint co-move with innovations to household labor income, both over a two-year interval as well as over a longer time horizon. The empirical question posed in this paper is not only important in its own right, but its answer is central for evaluating any forward-looking dynamic consumption-savings micro- or macroeconomic model in which households choose consumption and wealth accumulation in the presence of a stochastic labor earnings stream. In this paper, we argue that the short- and long-run consumption- and wealth responses contain useful information to identify the household consumption- and savings functions. For this purpose, the Italian Survey of Household Income and Wealth (SHIW) for the years 1991-2016 is ideal since it is one of the very few available household-level data sets that allows us to observe changes over time in all variables entering into the household budget restriction, and does so over a long time horizon for the same household.

In the theoretical consumption-savings models that our empirical analysis seeks to inform, the feasible consumption-savings choices of households crucially depend on the menu of financial and real assets available to them. Existing models differ starkly with respect to the assumptions regarding this menu. At one extreme, in so-called hand-to-mouth consumer models financial assets are entirely absent and consumption bears all the adjustment to income shocks. At the other extreme, the complete markets model (the underlying abstraction of any representative agent macro model) envisions a full set of state-contingent assets that households can trade without binding short-sale constraints. In this model wealth bears all the adjustment to an income shock, and consumption bears none (unless the income shock is an aggregate shock and cannot be diversified internationally). Our empirical findings on consumption and wealth adjustments are therefore informative about the insurance possibilities available to households. Distinguishing between these sets of models is not only important for positive questions (e.g., what is the joint income-consumption dynamics, the response of the macro economy to shocks, the pricing of financial assets, to name a few) but also for normative policy analysis. The desirability of social insurance policies (e.g., unemployment insurance, a redistributive tax code) depends crucially on how well households can privately (self-) insure against idiosyncratic income shocks, which in turn is determined by their access to and the sophistication of available asset markets.

Given the central importance of the question, it is perhaps not surprising that a sizable literature exists on this topic that we will review in the next section. However, most authors have focused on the consumption response to income shocks alone, but have not explicitly analyzed the corresponding response of the other components of the budget constraint (and specifically, the different forms of wealth) to the same shock. This can be mainly attributed to the scarcity of suitable panel data that contains repeated observations on both income on one hand, and consumption as well as wealth on the other hand, for the same set of households. In addition to performing this analysis, we also exploit the long panel dimension on income, consumption, and wealth in the SHIW to document the consumption and wealth response over long time horizons, concretely, up to six years. We then use this longer time horizon evidence to evaluate different partial insurance frameworks, starting from the original permanent income hypothesis and a structure that permits partial insurance even against permanent income shocks.

This paper seeks to make two broad contributions. The first is descriptive-empirical and documents a set of stylized facts concerning the co-movement of unexpected income changes and various measures of household consumption and wealth. The second contribution is to analyze the implications of these stylized facts for the properties of the household consumption- and savings (wealth accumulation) functions, from the perspective of the classical permanent income hypothesis (PIH) and simple extensions.

To do so, the analysis proceeds in four steps. In the first step, we construct a measure of idiosyncratic labor income shocks by regressing after-tax labor income on a set of observable household characteristics, which include age, education, age-education interactions, and time fixed effects. Our measure of income shocks is the change in the residuals from this regression. We then document co-movements, at the household level, of these labor income shocks with changes in other components of the household budget constraint, such as income from other sources, various consumption expenditure variables, and different measures of household wealth. In order to focus on households that indeed face labor income risk, we restrict our analysis to a sample of households whose head is between the age of 25 and 55 and is not retired. Our first finding (see Tables 2 and 3 and Figures 3, 4, and 5) is that consumption expenditures display only a modest co-movement (between 10 and 20 cents to the Euro) with income shocks, the other sources of income, including transfers, show negligible co-movement (less than 5 cents to the Euro) with income shocks, and wealth shows a large (exceeding 100 cents to the Euro) co-movement with income shocks.

The large wealth responses suggest that it is useful to divide the sample into two groups: households that do not own businesses or real estate (including their primary residence) and

households that do. Households in the first group comprise approximately 15% of the total sample. We find that for households in this group, non-durable consumption changes by about 35 cents in response to a short-run (two years) 1-Euro change in after-tax labor income, while financial wealth responds by about 30 cents (see Table 4). We also find that in response to longer-run (six years) income changes the consumption response becomes stronger, and the wealth response mildly increases as well with the time horizon. For the second group of households (those with businesses and/or real estate), we find that the consumption response to income shocks is significantly smaller while the wealth response is considerably larger. We therefore devote the last part of the paper to a more detailed empirical analysis of these households with housing and business wealth.

In the second step, we turn to theory and study what the consumption and wealth response, over short and over longer horizons, can teach us about the properties of the consumption- and the asset accumulation (saving) function. Concretely, we first assess whether the simplest variant of a consumption-savings model, a formalized version of the permanent income hypothesis in which infinitely-lived households¹ can freely borrow and save with a risk-free bond whose real return equals the subjective household time discount rate, face no binding borrowing constraints, have quadratic utility, and face both purely transitory and purely permanent shocks, can account for the empirical findings. In that model the consumption function and the saving function are available in closed form, and one can derive the consumption and wealth responses to an income shock analytically. We show that they are simple functions of the ratio between the variance of the permanent and the transitory shock, as well as the share of the transitory shock that is due to measurement error in income. We assess whether the co-movement between income, consumption, and wealth changes both in the short run and in the long run predicted by the PIH is quantitatively consistent with that observed in the data. We find that for our first sample (households without business and real-estate wealth), if permanent shocks are an important source of income risk and measurement error in income is modest, then the PIH captures the short-run consumption and wealth responses well. Quantitatively, and over a longer horizon, the PIH model predicts too strong a consumption response and too weak a wealth response, however, suggesting consumption- and savings functions that feature some (at least temporary) consumption insurance even against permanent shocks.

In the third step, motivated by the previous results, we extend the PIH towards a reduced-form partial insurance model in which even permanent shocks are partially insurable, consistent with the evidence in Blundell et al. (2008). We parameterize the degree

¹ We probe the importance of finite lives in Section 4.2 of the paper.

of such partial insurance and show that if it is sufficiently small (in the sense of a precise threshold), then the wealth response continues to *decline* with the time horizon N of the income shock as in the PIH (and the consumption response continues to increase with N). In contrast, if partial consumption insurance against permanent shocks is sufficiently large, then the wealth response rises in N , providing a qualitative distinction between the PIH and modest deviations from it on the one hand, and a model with more substantial consumption insurance against permanent shocks on the other hand. This result also demonstrates the broader point we want to emphasize in this paper, that observing longer-run responses of consumption and wealth to income shocks contains valuable information about the consumption function that can help distinguishing between different consumption-savings theories. We then use the observed longer-run wealth responses to income shocks to determine the degree of partial insurance (together with the degree of measurement error and the persistence of income shocks as above). Since wealth responses are moderately increasing with the time horizon, we find that a positive and substantial degree of insurance in which only about 2/3 of a permanent shock transmits into current consumption (and the rest into wealth) best fits the data.² Thus we conclude that a simple departure from the PIH that permits some limited insurance even against permanent shocks best describes our Italian household consumption- and wealth data, for the (highly selected) group of households that derive no income from either owned homes or businesses.

In the fourth and final step, in Section 5 we analyze the wealth response to income shocks for households who own real estate and/or business wealth. We document that real estate and business wealth co-moves especially strongly with labor income shocks. We argue that a large part of this co-movement may be driven by a correlation between labor income shocks and the *prices* of real estate and the *value* of businesses, rather than represent *wealth accumulation behavior* of households in response to labor income shocks. This leads us to conclude that the dominant set-up used to study heterogeneous agents economies (such as Aiyagari (1994), or Krueger et al. (2016), or Kaplan et al. (2018)), in which households only face idiosyncratic income shocks might be missing important sources of risk when it

² Although we are agnostic about the sources of this partial insurance by specifying a reduced-form consumption function with partial consumption insurance (which nests the PIH as a special case and point of comparison), we think of precautionary savings/buffer stock savings behavior as a primary candidate. Note that in the PIH version of the model there is no precautionary saving at all. In a model with CARA utility and absent borrowing constraints households engage in precautionary saving, but the amount they save for precautionary reasons is independent of their income or wealth level, and the realization of their income shock. Thus the PIH and the CARA utility versions of the incomplete markets consumption-savings model have exactly the same predictions for the consumption response to an income shock (and thus exactly the same predictions for the regression coefficients we estimate empirically). For a full theoretical treatment of the CARA case see Caballero (1990) and Wang (2003).

is applied to the entire population of households. This conclusion in turn motivates our sample selection in the parts of the paper where we evaluate simple consumption-savings theories.³

In the next section, we place our contribution into the existing empirical and theoretical quantitative literature. The data we use as well as the descriptive empirical results we derive are discussed in Section 3. In Section 4, we discuss what the descriptive results from Section 3 can teach us about the consumption-savings functions implied by simple partial equilibrium versions of incomplete markets consumption-savings models (the formalized PIH and an extension that allows for partial insurance against permanent shocks). Section 5 presents further evidence on the importance of changes in the value of real estate and business wealth associated with labor income shocks, and Section 6 concludes. Further details about the data, sample selection as well as detailed derivations for and extensions of the theoretical models employed in Section 4 are in the Appendix.

2 Related Literature

This paper builds on the large literature that has used household level data sets to quantitatively evaluate or formally statistically test the empirical predictions of Friedman (1957) permanent income hypothesis, and more broadly, assesses the response of household consumption and wealth to idiosyncratic income shocks. In this literature, Hall and Mishkin (1982) and Altonji and Siow (1987) represent seminal early contributions, and the initial body of work is discussed comprehensively in Deaton (1992). Jappelli and Pistaferri (2010a) and Jappelli and Pistaferri (2017) are comprehensive surveys of the subsequent empirical literature, Commault (2022) and Colarieti et al. (2024) are recent important contributions, and Kaplan and Violante (2022) summarize the corresponding predictions of standard incomplete markets models concerning this question. How strongly consumption responds to income shocks of a given persistence is the central question of this literature.⁴

How strongly consumption responds to income shocks has also been estimated for the U.S. and a number of other countries in the context of tests of perfect consumption in-

³ For a recent analysis on the impact of the price of real estate in incomplete market economies see Berger et al. (2017), and for recent evidence on the importance of the change of price of assets for wealth dynamics, see Fagereng et al. (2019). The importance of distinguishing between business owners and non-business owners in studies of consumption and wealth accumulation behavior is also emphasized by Hurst et al. (2010), and Kerr et al. (2017) provide a general discussion of the distinctive attributes of business owners.

⁴ How strongly household consumption responds to *predictable* changes in income is the subject of a large literature on excess sensitivity. In contrast, the excess smoothness literature studies how strongly household consumption adjusts in response to permanent income shocks. See e.g., Luengo-Prado and Sorensen (2008), and again the survey in Jappelli and Pistaferri (2017), chapter 8.

insurance.⁵ These tests do not need to distinguish between expected income changes and income shocks, and between transitory and permanent shocks since all income fluctuations ought to be smoothed and all shocks are fully insured, according to the perfect consumption insurance hypothesis.

Dynarski and Gruber (1997) and Krueger and Perri (2005, 2006) take a more agnostic view about the underlying (partial insurance) model generating the data and present the correlation between income and consumption changes as a set of stylized facts that quantitative models ought to match. The spirit of our empirical analysis is similar to these studies. For Italy, in a sequence of papers Jappelli and Pistaferri (2000a,b, 2006, 2010b, 2011) and Jappelli et al. (2008) employ the SHIW data to study the dynamics of household income, and the latter three study the joint dynamics of household income and consumption.⁶

Blundell et al. (2008) construct a consumption and income panel by skilfully merging data from the CEX and the PSID, and use this panel to estimate the extent to which households can insure consumption against transitory and permanent income shocks. Kaplan and Violante (2011) evaluate whether a class of incomplete markets models can rationalize the empirical estimates for consumption insurance that Blundell et al. (2008) obtain. Related, Aaronson et al. (2012) investigate the consumption response to an increase in the real wage in the U.S. Similar to our study, they find that the adjustment in real wealth (vehicles in particular, in their case) is a crucial feature in their data, and they construct a model with consumer durables to account for these facts. In a closely related paper, Fella et al. (2020) use a fully specified partial-equilibrium consumption-savings model as well as indirect inference (in the same spirit as Guvenen and Smith (2014) in their study of income risk and partial consumption insurance) to argue that precautionary saving is important to match the wealth regression coefficients over a longer time horizon.⁷

Finally, our work and the papers cited so far focus on the role of consumption self-insurance through financial markets in the face of stochastic labor income shocks. A complementary literature studies the importance of various insurance mechanisms against idiosyncratic *wage* risk, most notably, adjustments in family labor supply. See Heathcote et al. (2014), Blundell et al. (2016, 2018) and also Guner et al. (2012) and Holter et al. (2019) for important recent contributions.

⁵ See e.g., Mace (1991), Cochrane (1991), or Schulhofer-Wohl (2011) for the U.S. and Townsend (1994) or Mazzocco and Saini (2012) for India.

⁶ See Padula (2004) for an additional empirical study that uses the SHIW data to study the consumption response to permanent income shocks, but with the focus on consumer durables.

⁷ The first version of this paper appeared in 2011, and Fella et al. (2020) wrote their paper in response to our empirical regression estimates of varying horizon N and argue that a structural model with precautionary saving, rather than the pure version of the PIH, fits the data better.

3 Evidence

We use the only *long panel* data set that, to the best of our knowledge, contains detailed information about household income, consumption, and wealth: the Italian Survey of Household Income and Wealth (SHIW) for the years 1991-2016.

3.1 The Survey of Household Income and Wealth

The Survey of Household Income and Wealth (henceforth, SHIW) is conducted by the Bank of Italy. The survey started in 1965, but before 1987 it did not contain any panel dimension and did not contain complete wealth and consumption data. From 1987 until 2016 on the SHIW has been conducted every two years (with the exception of the 1995 and 1998 surveys which were conducted 3 years apart) and it includes about 8000 households per year, chosen to be representative of the whole Italian population. It also has a panel structure, and a fraction of households in the sample is present in the survey for repeated years. This data set is valuable and unique for our purposes as it contains *long panel* information for many categories of income, consumption, and wealth for each household.⁸ The panel dimension of income is particularly helpful for assessing the nature (i.e., permanent or temporary) of income changes and shocks. The fact that the data contains, for the same household, panel information on income, consumption, and wealth is crucial for inferring how a given household adjusts its consumption in response to an income change of a given type, and which and how various components of wealth change in association with income fluctuations.⁹

Table A1 in the Appendix displays the total sample size of the data from 1991 on (the first year for which comprehensive wealth data are available and thus the starting point of the sample we use in our analysis) as well as the share of the households in each wave of the SHIW that were already present in previous waves. We observe that the panel dimension of the data set since 1991 is substantial and has grown over time, with the fraction of the 7,420 households in the 2016 wave already being present in previous waves exceeding 50%.

Since the focus of this project is on the effects of earnings changes for households who are active in the labor market, we define an observation as a household who is in the survey

⁸ Jappelli and Pistaferri (2010b) show that aggregate consumption and aggregate income from the SHIW display growth rates that are very similar to the corresponding NIPA figures, suggesting that the coverage of the survey is comprehensive. See also recent work by Checchi et al. (2023) showing that inequality measures in the SHIW are comparable to those in administrative data

⁹ The US Consumer Expenditure (CEX) Survey has a panel dimension but the fact that it is short (only two periods), that observation periods for income and consumption do not perfectly coincide (see Gervais and Klein (2010) for a treatment of this problem), and the fact that there is no panel dimension for wealth makes it of limited use for our purposes.

for at least two consecutive periods and whose head is between the age 25 and 55 and is not retired in both periods. This leaves us with a sample of 18,661 observations over the period 1991-2016.

3.2 Organization of the Data and Measurement

In order to organize our empirical findings, we place them into the context of a sequential budget constraint of a standard incomplete markets model in which the household can self-insure by buying and selling a limited set of assets. This budget constraint reads as:

$$c_{nt} + c_{dt} + a_{t+1} + e_{t+1} = y_t + p_t + a_t + e_t + T_t, \quad (1)$$

where c_{nt} and c_{dt} denote consumption expenditures on non-durables (including rent and imputed rent for owner-occupied housing) and durables, respectively. a_{t+1} and e_{t+1} denote the *values* of the net asset position of financial and real wealth at the end of period t , whereas y_t measures after-tax labor income, T_t denotes net private and public transfers, and p_t denotes asset income, including income from financial assets (i.e., interests and dividends) and income from real wealth (rental income), correspondingly. Financial wealth includes liquid assets such as stocks and bonds whereas real wealth includes three types of less liquid assets, i.e., real estate, ownership shares of unincorporated business, and valuables (i.e., precious metals, art, etc.).

The Italian data is rich enough that we can measure all these variables for our households in the sample.¹⁰ The first step of our empirical analysis is to control for differences in family size across households by expressing all variables in adult equivalent units by dividing each observation by the appropriate OECD equivalence scale.¹¹ Table 1 below reports some basic summary statistics for our sample, and three separate sub-samples (pre- and post-2006), divided roughly by the midpoint of the overall sample.

Since our main focus is on income changes that are idiosyncratic and unpredictable (that is, on idiosyncratic income *shocks*) we first attempt to purge the data from aggregate effects and predictable (based on observables) individual changes by regressing each variable on time dummies, on a quartic in the age of the head of the household, on education and regional dummies, and on age-education interaction dummies. Our empirical exercise is then carried out on the residuals from these first-stage regressions.

¹⁰ For the exact variable definitions in the SHIW, please see Appendix A.

¹¹ This procedure has a minor impacts on the results. For labor income y_t , for example, around 99% of the cross-sectional variation of equalized income growth is due to variation in the growth rate of raw income.

We then denote by $\Delta^N x = \frac{x_t - x_{t-N}}{N}$ the annualized difference between an equivalized variable x today and N periods ago and we obtain, setting $N = 2$ (with the exception of 1998 where we set $N = 3$):

$$\begin{aligned} & \Delta^2 c_{nt} + \Delta^2 c_{dt} + \Delta^2 a_{t+1} + \Delta^2 e_{t+1} \\ = & \Delta^2 y_t + \Delta^2 p_t + \Delta^2 T_t \\ & + \Delta^2 a_t + \Delta^2 e_t \end{aligned} \tag{2}$$

In the rest of the paper we refer to the residualized income changes $\Delta^2 y_t$ synonymously as “income shocks” or (unpredictable) income changes.

Table 1: SHIW sample summary statistics

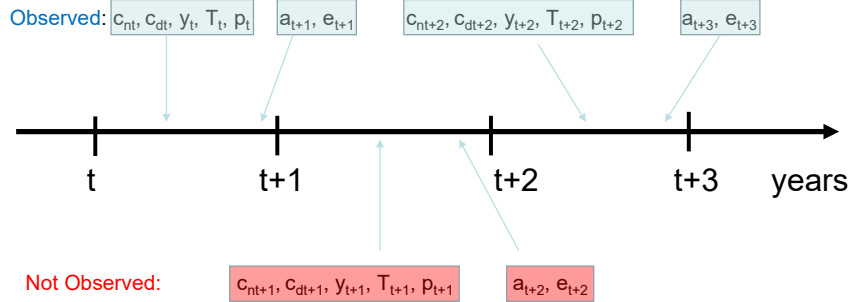
	Average Level			Annualized Growth		
	1991	2006	2016	1991-2006	2006-16	1991-2016
Age of head	40.3	42.1	44.5	0.3%	0.5%	0.4%
Household size	3.5	3.1	2.9	-0.9%	-0.5%	-0.7%
Labor income	9553	11192	9534	1.1%	-1.6%	-0.0%
Asset income	1924	2475	2183	1.7%	-1.2%	0.5%
Transfers	317	344	352	0.5%	0.2%	0.4%
Total consumption	9355	10529	9622	0.8%	-0.9%	0.1%
Non-durable consumption	8211	9526	8929	1.0%	-0.6%	0.3%
Durable consumption	1144	1003	693	-0.9%	-3.6%	-2.0%
Net financial wealth	6066	9246	7818	2.8%	-1.7%	1.0%
Net real estate wealth	37868	65058	54367	3.7%	-1.8%	1.5%
Net business wealth	6224	13555	6228	5.3%	-7.4%	+0.0%
Total net wealth	51737	89452	69862	3.7%	-2.4%	1.2%

Note: Summary statistics for the selected SHIW sample, for the years 1991, 2006, and 2016. All variables except age and household size are per adult equivalent and in 2000 Euros.

Note that, due to the biannual nature of our data set, the last two terms $\Delta^2 a_t$ and $\Delta^2 e_t$ cannot be observed in the data since wealth information is only available for the end of the period, and we observe every household only every two periods. This fact is clarified in Figure 1 which shows the frequency and exact timing with which different variables (including, crucially, wealth variables) are observed in the SHIW data set.

The empirical question we now want to answer is how the observable differences in the budget constraint co-move with changes in labor income $\Delta^2 y_t$.

Figure 1: Timeline in the Survey of Household Income and Wealth (SHIW)



Note: The diagram shows the available variables in the SHIW. Period t and $t + 2$ are survey years, whereas $t + 1$ is not.

3.3 Empirical Results


In Figure 2, we display the cumulative distribution function of observed residual annualized labor income changes (in levels and logs). The picture shows that a substantial fraction (about 10% of households) experience income changes that are larger than 2000 Euros (annualized, per adult equivalent) or larger than 20% of their labor income.

We fully acknowledge that a possibly significant share of this observed variation in labor earnings may be due to measurement error or to components that are predictable to the household but not to us, and thus will address these issues explicitly when comparing the stylized facts from the data to the predictions of the models we use to assess these facts.¹² To visualize the co-movement of various components of the budget constraint with income for each of the 20 bins of sorted income changes, we compute the average change in each observable component of the budget constraint and plot it against the corresponding income change. Figures 3-5 contain the results of this exercise, for non-durable and durable consumption, non-labor income components, and all forms of household wealth.

From Figure 3, we observe that non-durable consumption changes are positively correlated with income shocks. In addition, that relationship appears to be fairly linear. As we make precise below in Table 2, for the entire sample of households, on average a 1-Euro increase (decline) in after-tax labor income is associated with about a 9-cent increase (decline) in expenditures on non-durable consumption.

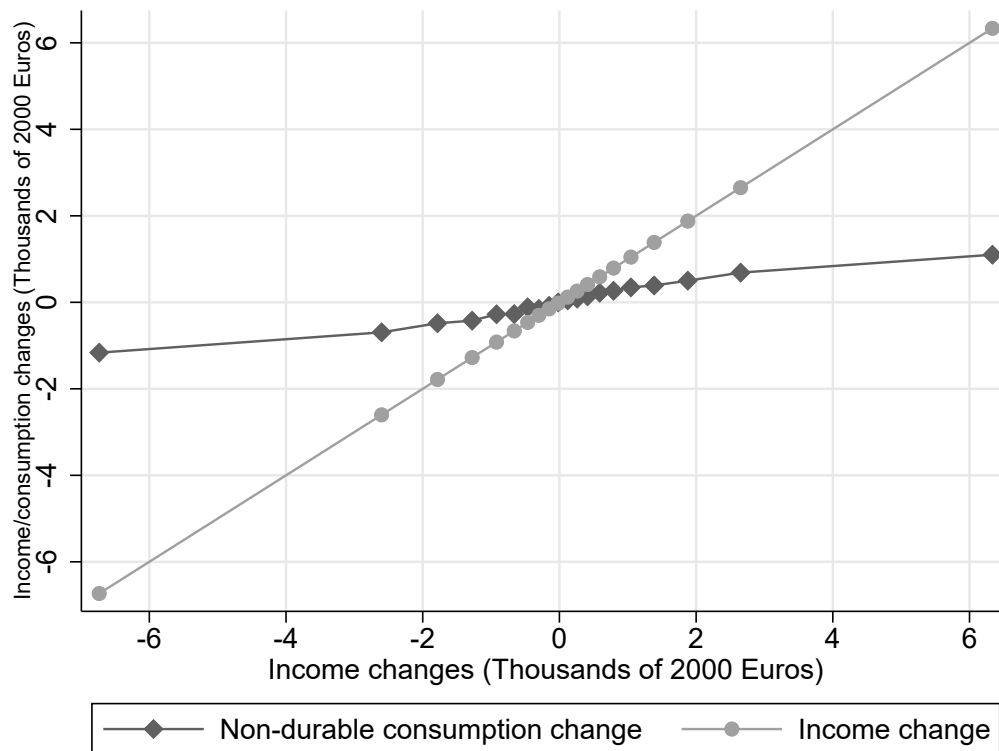
¹² Altonji and Siow (1987), in their critique of Hall and Mishkin (1982), stress the potential quantitative importance of measurement error in income changes or income growth for the type of regressions used in their and also in this paper. In Appendix C we also briefly discuss how the presence of income changes that are predictable to the household but not to us would change the interpretation of our results.

Figure 2: Cumulative distribution function of residual labor income changes




Note: The left panel is for changes in log of residual labor income, the right panel is for changes in levels of residual income. Labor income is after-tax wage income plus fringe benefits plus business income, per adult equivalent, in constant 2000 Euros. All changes are annualized.

Figure 3: Changes in labor income and consumption



Note: Households are sorted (by the size of the residual income change) into 20 bins, and the average residual consumption change and residual labor income change (in each bin) is plotted against the average residual labor income change. Each bin contains around 930 households. Labor income is after-tax wage income plus fringe benefits plus business income. Consumption is expenditures on non durable goods and services. All variables are per adult equivalent, in constant 2000 Euros. All changes are annualized.

Figure 4: Changes in labor income, durable consumption, transfers, property income, and financial income



Note: Households are sorted (by the size of the residual labor income change) into 20 bins, and the average residual variable change and residual labor income change (in each bin) is plotted against the average residual labor income change. Each bin contains around 930 households. Labor income is after-tax wage income plus fringe benefits plus business income. Transfers include private and public. All variables are per adult equivalent, in constant 2000 Euros. All changes are annualized.

In Figure 4, we display the co-movement of after-tax labor income with the other parts of household income, in particular transfer income (the upper right panel), and capital income from both real assets and financial assets (the lower two panels). The upper left panel shows the change in expenditures on consumer durables (mainly cars and furniture) for each income change bin. We observe that changes in expenditures on consumer durables co-move mildly positively with income shocks but less so than changes in expenditures on non-durables. Changes in labor income and in income from properties and from financial assets changes are, broadly speaking, uncorrelated with each other.¹³ On the other hand, there is a visible, significant, but quantitatively small negative co-movement between labor income changes and the change in net public and private transfers received by households. This negative correlation is especially noticeable for households with large income increases and income declines.

Figure 5 shows instead the co-movement of changes in various wealth components with labor income and shows how total wealth and all its components (financial wealth, real estate wealth, and business wealth) strongly co-move with labor income. This co-movement is most pronounced for business wealth and real estate wealth, and more modest for financial wealth.

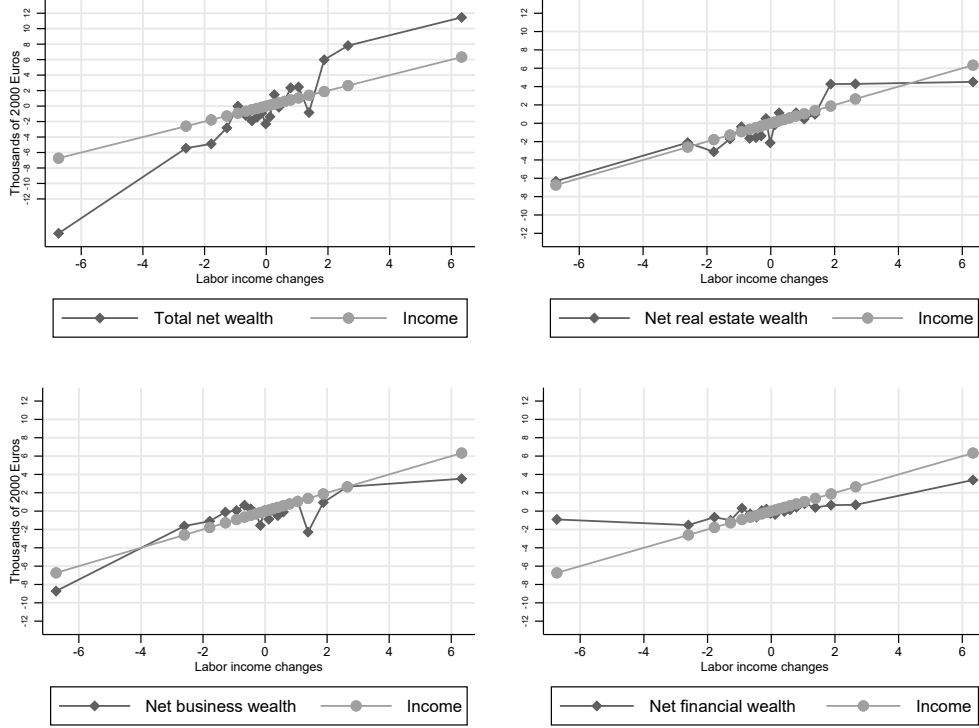
In order to formally evaluate the magnitude of the average response of the various components of the budget constraint to income changes we now run bi-variate regressions of the changes in the various components of the budget constraint on the changes in income. The results are reported in Tables 2 and 3 below. Since the OLS estimates, in particular for the wealth observations, may be influenced by a few large outliers that report large positive or negative changes in wealth, we also report the median regression (MR) estimates resulting from minimizing the sum of the absolute values of the residuals, rather than the sum of squared residuals. By putting less weights on extreme observations, MR estimates are more robust to the influence of outliers.

Results in Table 2 quantitatively confirm the visual evidence from Figures 3 and 4 that changes in expenditures on consumer non-durables Δc_n and on durables Δc_d are significantly associated with changes in income but are much smaller than the income changes. On average when income changes by 1 Euro total consumption expenditures change by about 11 cents.

The figures above also show that the other sources of income are only weakly correlated with labor income changes. Table 2 splits total net transfers T into transfers from family

¹³The sum of changes of income from properties and from financial assets corresponds to the term Δp in equation (2).

Figure 5: Changes in labor income and changes in different components of wealth



Note: Households are sorted (by the size of the residual labor income change) into 20 bins, and the average residual variable change and residual labor income change (in each bin) is plotted against the average residual labor income change. Each bin contains around 930 households. Labor income is after-tax wage income plus fringe benefits plus business income. All variables are per adult equivalent, in constant 2000 Euros. All changes are annualized.

Table 2: Co-movement of budget constraint components with changes in labor income

	Δc	Δc_n	Δc_d	ΔT	ΔT_F	ΔT_O	Δp
β_{OLS}	11.2	8.7	2.4	-1.8	-3.3	1.5	0.6
	(4.48)	(2.49)	(2.23)	(0.89)	(0.56)	(1.28)	(0.53)
R^2	0.03	0.04	0.00	0.01	0.01	0.00	0.00
β_{MR}	21.9	17.8	0.8	-0.3	-0.2	-0.1	1.0
	(0.26)	(0.22)	(0.03)	(0.02)	(0.01)	(0.02)	(0.05)
R^2	0.03	0.03	0.00	0.00	0.00	0.00	0.00
<i>Obs.</i>	18661	18661	18661	18661	13976	13976	18661

Note: Bivariate regression coefficients with income changes as the independent variable. (OLS or least absolute deviations). Standard errors (for OLS) are clustered at the household level and are in parentheses.

and friends T_F and other transfers T_O (which includes pensions and arrears) and indicates that the former accounts for the majority of the (not very large) negative correlation between labor income changes and changes in transfers.¹⁴ The adjustment of family transfers for a Euro in lower labor income is in the order of 3 cents. The existence and negative correlation between labor income changes and changes in family transfers may lend some qualitative support to models that permit household to engage in more explicit insurance arrangements than the simple self-insurance through asset trades that standard incomplete markets models envision (e.g., models with private information or limited commitment). Note, however, that these changes in transfers and their correlation with labor income changes are quantitatively very small. Finally, changes in asset income Δp are only very weakly correlated with income shocks, as the last column of Table 2 suggests.¹⁵

Table 3: Co-movement of wealth components with changes in labor income

	$\Delta(a + e)$	Δa	Δe^{re}	Δe^b	Δe^v
β_{OLS}	313.7	24.3	114.0	174.5	0.9
	(75.2)	(14.4)	(34.8)	(42.6)	(1.4)
R^2	0.05	0.01	0.02	0.04	0.00
β_{MR}	128.7	21.7	42.2	12.3	2.4
	(2.04)	(0.38)	(1.40)	(0.34)	(0.10)
R^2	0.01	0.01	0.00	0.00	0.00
<i>Obs.</i>	18661	18661	18661	18661	18661

Note: Bivariate regression coefficients with income changes as the independent variable. (OLS or least absolute deviations). Standard errors (for OLS) are clustered at the household level and are in parentheses.

Results in Table 3 confirm the findings from Figure 5 that changes in labor income are strongly associated with changes in wealth. The first column reports the result of regressing residual changes in total wealth on residual changes in labor income while the subsequent columns report the results using financial wealth (a), real estate wealth (e^{re}), business wealth (e^b), and valuables (e^v). Notice that results change quantitatively very significantly whether we use OLS or MR regressions, suggesting that there are some households reporting very large changes in wealth (in particular, business wealth) which strongly affect the OLS results.

¹⁴ Note that the lower number of observation in the T_F and T_O regression is due to the fact that disaggregated data on transfers are not available in the early survey years.

¹⁵The measure of capital income does not include capital gains, and therefore we acknowledge that if capital gains are positively correlated with our measure of idiosyncratic income shocks, then this estimate might potentially be downward biased.

The upshot of the table though is that, regardless of the regression method, on average a 1-Euro change in labor income is associated with changes in wealth that are larger than 1 Euro, and that this finding is mainly driven by movements in business wealth and real estate wealth. This result suggests that a simple consumption-savings model in which households are subject solely to income shocks cannot be consistent with this fact.¹⁶ We conjecture that the main reason for this result is the presence of shocks to the value of the wealth which are *correlated* with the value of labor income. An example of this would be an entrepreneur that receives a positive shock to the value of her business which at the same time raises both her measured labor income and her wealth. Another example would be a city-specific shock which raises, at the same time, labor income and real estate wealth of its residents.

Therefore, in order to isolate a household’s response to a “pure” income shock we now select households which report neither income from business nor from real estate.¹⁷ We think of these households as providing us with the most plausible set to evaluate the predictions of simple consumption-savings models. Since these models study the consumption-savings response to idiosyncratic income shocks, typically *after* the government has provided some income insurance through the tax-transfer system, we now add to our income measure (and thus our income shock measure) used thus far government transfers, and document the consumption- and wealth response of the selected sub-sample to this measure of disposable income.¹⁸

The key result to notice from Table 4 is that for this selected sample non-durable consumption co-moves significantly more strongly and wealth co-moves significantly less with income. The non-durable consumption response is in the order of 35 cents for each Euro, and the response of wealth amounts to approximately 30 cents. In the next section, we assess whether, as a first basic check of consumption theory, the standard formalized version of the permanent income hypothesis in the spirit of Friedman (1957) provides a reasonable approximation of the data for this selected group of households. This analysis also provides some guidance along what dimension this basic model ought to be extended to match the co-movement facts for the whole sample of households.

¹⁶ Note that this large change in the real value of assets is not in principle inconsistent with the budget constraint. If income in period $t - 1$ (which we do not observe, due to the biannual structure of the data set) were highly correlated with income change $y_t - y_{t-2}$ then the right-hand side of the budget constraint could change by more than 1 Euro for each Euro in Δy . In practice though, for empirically relevant income processes this correlation is not high enough to generate such a large response of wealth.

¹⁷ Income from real estate includes imputed rents on owned properties, therefore our sample exclude all homeowners.

¹⁸ In the descriptive analysis thus far, in contrast, the objective was to document which items of the household budget constraint, including private but also public transfers, co-move with labor income after taxes, and how strong that co-movement is.

Table 4: Co-movements for selected sample

	Δc_n	Δc_d	Δa
β_{OLS}	34.7 (3.2)	1.4 (2.3)	29.7 (4.4)
R^2	0.15	0.00	0.04
β_{MR}	31.4 (1.3)	1.1 (0.2)	15.7 (1.5)
R^2	0.08	0.00	0.02
<i>Obs.</i>	2612	2612	2612

Note: Bivariate regression coefficients with income changes as the independent variable, for sample of households without real estate or business income. (OLS or least absolute deviations). Standard errors (for OLS) are clustered at the household level and are in parentheses.

4 Theory

4.1 The Permanent Income Hypothesis

We now want to investigate whether versions of a standard incomplete markets model are consistent with the facts displayed in the previous section. In this section, we summarize the empirical predictions of a model based on the permanent income hypothesis for the question at hand, and evaluate to what extent the empirical evidence presented above is consistent with this model. In the next section, we then study a calibrated version of a standard incomplete markets life cycle model with a precautionary savings motive.

Suppose that infinitely lived households have a quadratic period utility function, can freely borrow and lend¹⁹ at a fixed interest rate r , discount the future at time discount factor β that satisfies²⁰ $\beta(1+r) = 1$, and face a process of measured *after-tax* labor income of the form

$$\tilde{y}_t = y_t + \gamma_t = \bar{y} + z_t + \varepsilon_t + \gamma_t \tag{3}$$

$$z_t = z_{t-1} + \eta_t \tag{4}$$

where \tilde{y}_t is measured (by the econometrician) labor income, $\gamma_t \sim N(0, \sigma_\gamma^2)$ is classical measurement error in income, y_t is true labor income that enters the household's budget

¹⁹ Of course a no-Ponzi condition is required to make the household decision problem have a solution.

²⁰ The predictions of the theory we focus on below, especially how the consumption- and asset responses vary with the horizon over which the income shocks are observed are robust to assuming $\beta(1+r) < 1$. However, if $\beta(1+r) > 1$, then households would accumulate wealth indefinitely with an infinite planning horizon, and our ensuing analysis no longer applies (but of course there could never be a stationary equilibrium with such an interest rate either.)

constraint, \bar{y} is expected household income, $\varepsilon_t \sim N(0, \sigma_\varepsilon^2)$ is a transitory income shock, and $\eta_t \sim N(0, \sigma_\eta^2)$ is a permanent income shock. The shocks $(\varepsilon_t, \eta_t, \gamma_t)$ are assumed to be uncorrelated over time and across each other.

Aggregating across wealth components and focusing on non-durable consumption, the household faces a budget constraint of the form

$$c_t + w_{t+1} = y_t + (1 + r)w_t \quad (5)$$

where $w_t = a_t + e_t$ is total wealth and c_t are expenditures on non-durable consumption, including (imputed) rent for housing. We show in the Appendix how a model that includes housing explicitly can be reduced to the formulation studied in this section as long as there are competitive rental markets, and the stock of housing can be adjusted without any frictions or binding financing constraints. In addition, for the empirical implementation of this model we include transfers T_t as part of after-tax labor income.

4.1.1 Empirical Predictions

Under the maintained assumptions, the household optimally holds consumption constant in expectation, $E_t c_{t+1} = c_t$, and the optimal consumption function reads as

$$c_t = r w_t + \frac{r}{1+r} E_t \sum_{\tau=0}^{\infty} \frac{y_{t+\tau}}{(1+r)^\tau} = r w_t + \bar{y} + z_{t-1} + \eta_t + \frac{r}{1+r} \varepsilon_t. \quad (6)$$

As is well-known, the realized changes in income, consumption, and wealth in this model are then given by (see e.g., Deaton (1992)):

$$\begin{aligned} \Delta c_t &= \frac{r}{1+r} \varepsilon_t + \eta_t \\ \Delta w_{t+1} &= \frac{\varepsilon_t}{1+r} \\ \Delta \tilde{y}_t &= \eta_t + \Delta \varepsilon_t + \Delta \gamma_t \end{aligned} \quad (7)$$

where $\Delta x_t = x_t - x_{t-1}$.

Equipped with these results, we can now deduce the consumption and wealth responses to income changes, as measured by the same bi-variate regressions we ran for our Italian data. First, since we have available a full panel and the survey is carried out only two periods, we need to work with changes of variables over N periods, which are given by:

$$\Delta^N x_t = x_t - x_{t-N} = \Delta x_t + \Delta x_{t-1} + \dots + \Delta x_{t-N+1}.$$

Using (7), we find that

$$\begin{aligned}
\Delta^N c_t &= \sum_{\tau=t-N+1}^t \left(\frac{r\varepsilon_\tau}{1+r} + \eta_\tau \right) \\
\Delta^N w_{t+1} &= \sum_{\tau=t-N+1}^t \frac{\varepsilon_\tau}{1+r} \\
\Delta^N \tilde{y}_t &= \sum_{\tau=t-N+1}^t \eta_\tau + \Delta^N \varepsilon_t + \Delta^N \gamma_t
\end{aligned} \tag{8}$$

and thus the bi-variate regression coefficients of N -period consumption and wealth changes on N -period income change are given as

$$\begin{aligned}
\beta_c^N &= \frac{Cov(\Delta^N c_t, \Delta^N \tilde{y}_t)}{Var(\Delta^N \tilde{y}_t)} = \frac{Cov\left(\sum_{\tau=t-N+1}^t \left(\frac{r\varepsilon_\tau}{1+r} + \eta_\tau\right), \sum_{\tau=t-N+1}^t \eta_\tau + \Delta^N \varepsilon_t + \Delta^N \gamma_t\right)}{Var\left(\sum_{\tau=t-N+1}^t \eta_\tau + \Delta^N \varepsilon_t + \Delta^N \gamma_t\right)} \\
&= \frac{N\sigma_\eta^2 + r\sigma_\varepsilon^2/(1+r)}{N\sigma_\eta^2 + 2(\sigma_\varepsilon^2 + \sigma_\gamma^2)} \\
\beta_w^N &= \frac{Cov(\Delta^N w_t, \Delta^N \tilde{y}_t)}{Var(\Delta^N \tilde{y}_t)} = \frac{Cov\left(\sum_{\tau=t-N+1}^t \frac{\varepsilon_\tau}{1+r}, \sum_{\tau=t-N+1}^t \eta_\tau + \Delta^N \varepsilon_t + \Delta^N \gamma_t\right)}{Var\left(\sum_{\tau=t-N+1}^t \eta_\tau + \Delta^N \varepsilon_t + \Delta^N \gamma_t\right)} \\
&= \frac{\sigma_\varepsilon^2}{(1+r)[N\sigma_\eta^2 + 2(\sigma_\varepsilon^2 + \sigma_\gamma^2)]}.
\end{aligned}$$

Conditional on a real interest rate r , these regression coefficients can be expressed exclusively as functions of the ratio of the size of permanent to transitory shocks, $Q = \frac{\sigma_\eta^2}{\sigma_\varepsilon^2 + \sigma_\gamma^2}$, and the share of transitory income shocks attributed to measurement error, $M = \frac{\sigma_\gamma^2}{\sigma_\varepsilon^2 + \sigma_\gamma^2}$.²¹

²¹ The estimated coefficient β_c^N can be decomposed into the regression coefficient obtained if income was measured without error, β , and the attenuation bias stemming from measurement error:

$$\beta_c^N = \beta \times \frac{1}{1 + \frac{2\sigma_\gamma^2}{N\sigma_\eta^2 + 2\sigma_\varepsilon^2}}$$

where

$$\begin{aligned}
\beta &= \frac{Cov\left(\sum_{\tau=t-N+1}^t \left(\frac{r\varepsilon_\tau}{1+r} + \eta_\tau\right), \sum_{\tau=t-N+1}^t \eta_\tau + \Delta^N \varepsilon_t\right)}{Var\left(\sum_{\tau=t-N+1}^t \eta_\tau + \Delta^N \varepsilon_t\right)} \\
&= \frac{N\sigma_\eta^2 + r\sigma_\varepsilon^2/(1+r)}{N\sigma_\eta^2 + 2\sigma_\varepsilon^2}
\end{aligned}$$

Using these definitions, we find

$$\beta_c^N = \frac{NQ + (1 - M)\frac{r}{1+r}}{NQ + 2} \quad (9)$$

$$\beta_w^N = \frac{1 - M}{(1 + r)[NQ + 2]}. \quad (10)$$

Straightforwardly, the larger is the size of the permanent shock, relative to the transitory shock, as measured by Q , the larger is the consumption response β_c^N and the smaller is the wealth response β_w^N . Second, increasing the period length N acts exactly like an increase in Q (notice that N and Q appear in the expressions above as a product exclusively). Transitory shocks are mean-reverting of the horizon of N years, whereas all permanent shocks during the N year accumulate in income changes, see equation (21). Therefore, an increase in N effectively increases the persistence of income shocks, and thus the PIH implies that the coefficient β_c^N is increasing in N and β_w^N is decreasing in N . To evaluate this last prediction in particular requires panel data on labor income, consumption, and wealth, which the Italian data, uniquely among household level data sets for industrialized countries, provides.

Larger measurement error lowers both coefficients due to the standard attenuation bias: it increases the variance of observed income, but leaves consumption and wealth unaffected since it is only income variation observed by the econometrician, but not experienced by the household. From equation (9), we observe that the share of measurement error is quantitatively unimportant for β_c^N for plausible values of r . True transitory shocks to income translate into consumption with a factor $\frac{r}{1+r} \approx 0$, while measurement error has an impact of exactly 0. Thus, to a first approximation, the share M of measurement error does not affect β_c^N . On the other hand, true transitory income shocks translate into changes in wealth one for one, whereas measurement error does not have any impact on the changes in wealth. Therefore, the degree of measurement error M has a strong impact on β_w^N , as (10) shows.

Finally, we observe that the *size* of the income innovations, σ_ε^2 and σ_η^2 , per se has no

so that

$$\begin{aligned} \beta_c^N &= \frac{N\sigma_\eta^2 + r\sigma_\varepsilon^2/(1+r)}{N\sigma_\eta^2 + 2\sigma_\varepsilon^2} \times \frac{1}{1 + \frac{2\sigma_\gamma^2}{N\sigma_\eta^2 + 2\sigma_\varepsilon^2}} \\ &= \frac{N\sigma_\eta^2 + r\sigma_\varepsilon^2/(1+r)}{N\sigma_\eta^2 + 2\sigma_\varepsilon^2 + 2\sigma_\gamma^2} \end{aligned}$$

We observe how the size of the bias in β_c^N is decreasing in N and Q . Thus another useful aspect of the longer panel dimension of the Italian data set is that it allows us to use income changes over longer time periods which mitigates the problem of (classical) measurement error in income.

impact on the regression coefficients. This is to be expected since quadratic utility and the absence of binding borrowing constraints implies that the household consumption and wealth choices obey certainty equivalence, and a precautionary savings motive is absent. In the next subsection, we will evaluate how important the incorporation of a precautionary savings motive is to rationalize the empirically observed co-movement of labor income, consumption, and wealth.

4.1.2 Evaluating the Empirical Predictions

We now ask whether for the sample of households that we identified in the empirical section as most appropriately modeled by the PIH, households without business and real estate wealth, the PIH is consistent with data. First, we let $N = 2$ and look at the minimal panel dimension, which in turn contains the maximal number of households in the data. For concreteness, we assume a real interest rate of $r = 0\%$ for the rest of this section.²² Equations (9)-(10) show that the exact value of the real interest rate affects the predicted values for (β_c^2, β_w^2) only insignificantly. We then ask what values of Q and M are needed to assure that the model predicts the same regression coefficients as in the data.

Recall that the empirical regression results for the sub-sample under question delivered a consumption response of $\beta_c^2 = 0.347$ and a financial wealth response of $\beta_w^2 = 0.297$. Using equations (9)-(10) we can determine which degree of income persistence Q and measurement error M is required for the model to match the data perfectly along these two stylized facts.²³ The results are $Q = 0.53$ and $M = 0.09$. As the discussion above indicates, the empirical consumption response of 34.7 cents for each Euro implies, for the PIH to be consistent with this fact, that income shocks are to a significant extent driven both by transitory shocks and by permanent shocks (since permanent shocks imply a one-for-one consumption response and transitory income shocks a (close to) zero consumption response). As discussed above, the size of measurement error plays essentially no role (and if $r = 0$, no role at all) for the consumption regression coefficient in the model. Conditional on a value for Q determined from the consumption data, the empirical wealth response then determines the required

²² As equations (9) and (10) make clear, the concrete value for the interest rate is quantitatively unimportant as long as it is close to zero. Furthermore, this assumption is empirically not unreasonable for Italian households using safe assets (such as bank accounts) during the sample period.

²³ Given equations (9)-(10), we can simply solve for Q and M given the observed β_c^2, β_w^2 as

$$Q = \frac{\beta_c^2 - r\beta_w^2}{1 - \beta_c^2 + r\beta_w^2}$$

$$M = 1 - \frac{2(1+r)\beta_w^2}{1 - \beta_c^2 + r\beta_w^2}$$

Table 5: Results for Longer N

N	No. Obs.	β_c^N		β_w^N	
		Data	PIH	Data	PIH
2	2612	34.7 (2.8)	34.7	29.7 (4.4)	29.7
4	1256	33.1 (3.5)	51.5	31.7 (12.4)	22.1
6	614	41.2 (4.7)	61.5	31.6 (8.6)	17.5

Note: Bivariate regression coefficients for different horizons N , with income changes as the independent variable and nondurable consumption and wealth as dependent variables. Data (with standard errors clustered at the household level in parentheses) and PIH model, as implied by equations (9) and (10).

degree of measurement error, and since the wealth regression coefficient is sizable as well, the PIH rationalizes this with relatively small measurement error in income.

With a choice of $Q = 0.53$ and $M = 0.09$ the PIH model matches the consumption and financial wealth response to labor income shocks over a two-year horizon by construction. Thus this fact cannot be interpreted as a success of the model per se. However the inferred value for Q can be validated using the panel dimension for labor income data, which, conditional on the particular form of the income process, can be used to derive an estimate for Q that only uses income data. Following the procedure outlined in Jappelli and Pistaferri (2006), and using only income data for the restricted sample of households without income from business or from real estate, we estimate $Q = 0.60$, which is close the value for $Q = 0.53$ inferred using consumption and wealth data.²⁴

Before turning to the precautionary savings model we now more fully exploit the unique panel dimension of the Italian data to evaluate the predictions of the PIH for income shocks over longer time horizons, that is, for increasing N . An increase in N means that more permanent shocks have accumulated, and that consumption should respond more strongly to a given income change. In Table 5, we summarize how the model-implied consumption regression coefficients vary with N . Since the sample size falls significantly as N increases, we restrict attention to $N \leq 6$. The model results are derived under the assumptions that $M = 0.09$ and $Q = 0.53$, the values needed for the model to exactly match the data for $N = 2$ and wealth being interpreted as financial wealth.

We observe that, as discussed earlier, the model predicts the expected increase in the consumption coefficients, but implies a decline in the wealth coefficients with the time

²⁴See Appendix B for the details of the estimation

horizon N . For consumption, the data suggests the same qualitative pattern, although the increase in the data is noticeably smaller than implied by the model. In contrast, the pattern of the financial wealth response to income shocks is qualitatively inconsistent with the data which display a mild increase in the wealth response as the time horizon increases from $N = 2$ years to $N = 6$ years. Note that the findings for $N = 4, 6$ provide a true test for the model as all model parameters have been chosen only with the data for $N = 2$ serving as targets.

To summarize, we conclude that the simple PIH model is successful in reproducing the empirically observed dynamic consumption response to income shocks of various duration. There are, however, three empirical observations that this model has trouble in rationalizing. First, the required degree of measurement error to match the short-run consumption- and wealth response is close zero, less than what is plausible given the wealth of evidence on the pervasiveness of measurement error in the data. Second, in contrast to the qualitative prediction of the model that income shocks should less strongly transmit into wealth as the horizon N increases, the income-wealth correlation display the opposite pattern. In the next section we evaluate whether introducing a simple model of partial consumption insurance against even permanent income shocks allows for a more plausible value of measurement error and delivers wealth regression coefficients that are mildly increasing rather than (strongly) decreasing with the time horizon N , as the pure PIH implies.

Finally, the PIH cannot match the observed large income-wealth correlations if wealth is interpreted more broadly to include real estate wealth (and business wealth), an interpretation that is mandated by a model that includes real estate explicitly (see Appendix F). We therefore, in Section 5, investigate further what could explain the observed large positive correlation between labor income shocks and real estate and business wealth.

4.2 Key Assumptions and Robustness

In this section we discuss the key theoretical assumptions that we have imposed in the previous section and that allowed us to derive the empirical predictions for the consumption- and wealth regression coefficients in closed form.²⁵ First, after-tax income is the sum of a permanent and a transitory component (where the latter could include measurement error). This has two main implications. First, and independent of the behavioral model of consumption, longer-run income changes $\Delta^N \tilde{y}_t$ are dominated by the accumulation of

²⁵ We thank our discussant Martin Holm for important comments on our work that have led to this section. See Holm (2023) for the core arguments we address here.

permanent shocks, as equation (21) shows.²⁶ Thus, the larger is N the more informative are the regression coefficients about responses to permanent shocks, and the less of a concern is classical measurement error in consumption that contaminates true transitory income shocks (see footnote 21 above). Second, and conditional on the consumption-savings behavior governed by the PIH being true, the consumption response to income shocks (over shorter or longer horizon) is directly informative about the share of the additive shocks that is permanent, as equation (9) clarifies. In fact, for $r = 0$, for any time horizon N there is a one-for-one relationship between the OLS regression coefficient and the share of income shocks that is permanent.

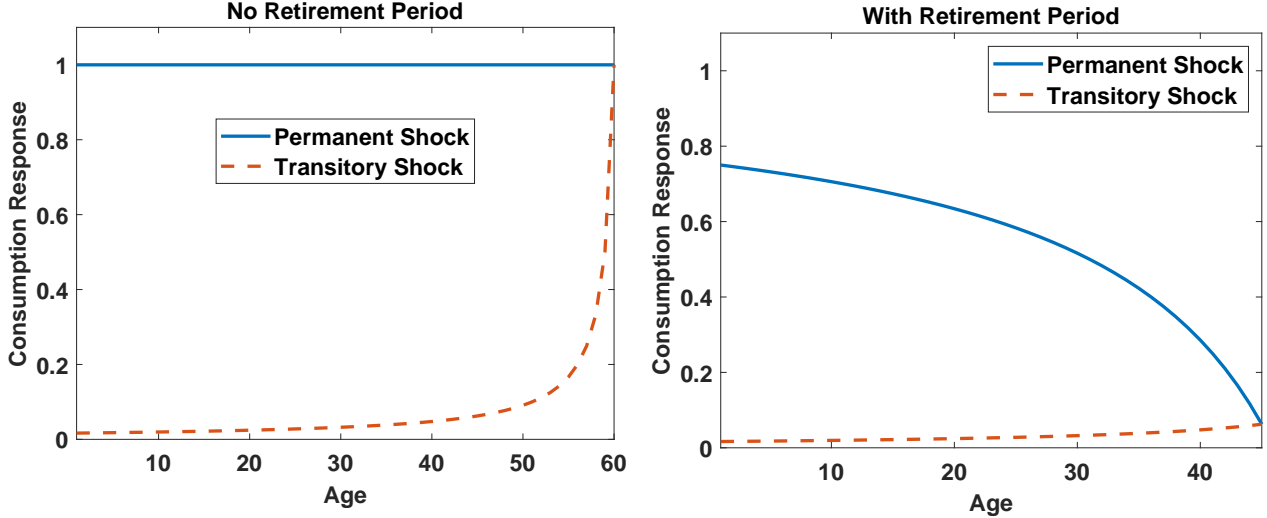
This latter result (which uses the consumption response, or, alternatively, the savings response, to a given income shock), highlights the second key assumption made in the previous section, namely that households have an infinite planning horizon. In Appendix D we discuss the consumption response to a permanent shock η_t and a transitory shock ε_t when individuals live until age $T < \infty$ and receive after tax labor income until retirement age $J \leq T$. Figure 6 displays these responses.

The left panel shows the consumption response to a permanent and a transitory shock received at a specific age, under the assumption that the last period of work J is also the last period of life, T . We make two observations: first, towards the end of life there is no difference in the consumption response to a transitory and a permanent shock, or in other words, towards the end of life all income shocks are permanent and consumption responds one-for-one to both shocks. Second, as long as individuals are sufficiently young (the remaining lifetime horizon is at least 30 years), the implication of the pure PIH and its finite lifetime version have quantitatively the same implications. The response to a permanent shock remains at 1, and the response to a transitory shock early in life is very close to zero (it is exactly zero if the lifetime horizon approaches ∞).

The right panel shows that if there is a period of prolonged retirement (and if pension income does not depend on the income shocks during working life as is the case in a pure Beveridgean pension system), these two main conclusions remain intact, but with a qualification. As before, towards the end of working life the response to permanent and to transitory shocks converges, but now to a magnitude that is closer to that of a transitory

²⁶ We should note that, as is common in the PIH literature with quadratic utility, we use a permanent-transitory decomposition in income *levels*, which allows for analytical expressions of the OLS regression coefficients implied by the theory. Holm (2023), in contrast, uses a similar decomposition for log-income, as is common in the literature on precautionary saving models with CRRA utility, although the basic insight that over longer horizon income changes are dominated by permanent shocks is common to both specifications, there is one key difference: with log-income, permanent shocks change the extent of future income risk (since income shocks become multiplicative) and thus the magnitude of the precautionary saving motive, which is absent in the pure form of the PIH discussed in the previous section.

Figure 6: Consumption response to permanent and transitory income shocks with finite lifetime



Note: The left panel assumes that the end of working life (age 60) corresponds with the end of life. The right panel assumes that the end of working life is age 45, and the end of life is age 60 (i.e., 15 years of retirement) and that the income shock does not affect pension benefits. We assume $r = 0$ as before.

shock (since the shock during the last working period is smoothed over the retirement period). Second, for young households the clear distinction between transitory and permanent shocks resurfaces and the predictions approach those of the previous section; however, since the permanent shock is not truly permanent (it ceases to have relevance at retirement), the consumption response remains noticeably below one.

Overall, this discussion suggests that if the assumption of finite life is important, it is so for individuals close to retirement. In order to probe the robustness of our empirical findings, in Table 6 we display our consumption and wealth regressions from Table 4 (and the first row of Table 5), but now for two sub-samples of roughly equal size, one composed of young households (households with a head aged 40 or younger) and one consisting of older households (those age 47 or older).²⁷ Comparing the estimates from the entire sample in Table 4 to those of the younger sample (first row of Table 6), we see that although the consumption response is somewhat larger and the wealth response a little smaller for the younger sample, our benchmark results are robust to the use of only younger households.

Interestingly, for the older sample (ages 47 to 55) the consumption response is smaller and the wealth response is larger than for the young (or the overall) sample. Of course, there

²⁷ Recall that the original sample already restricted attention to households between the ages of 25 and 55.

Table 6: Results for Young and Older Sample

	No. Obs.	β_c^2	β_w^2
Younger, $age \leq 40$	960	36.8 (4.1)	25.7 (6.5)
Older, $age \geq 47$	979	25.5 (5.7)	33.2 (6.9)

Note: Bivariate regression coefficients for young and old sample, with income changes as the independent variable and nondurable consumption and wealth as dependent variables. Standard errors clustered at the household level in parentheses.

are many possible explanations for this observation (including the possibility that borrowing constraints are more relevant for younger households, as a large empirical literature has argued), in light of Figure 6 one explanation is that for older households more of the income shocks are transitory either because of the nature of the shock (for young households the blue solid line in the left panel is more relevant, for older households the dashed red line), or because for older households who expect to live well beyond age 55, the permanent income shocks are not all that permanent (because retirement is close, see the solid blue line in the right panel of Figure 6).

4.3 Partial Insurance against Permanent Income Shocks

The pure form of the permanent income hypothesis implies certainty equivalence. There is no scope for precautionary saving, and consumption responds one-for-one to permanent income shocks, with financial wealth being unaffected. This was the basis for the observation in the last section that over longer horizons consumption co-varies more strongly with income, and wealth co-varies less strongly with labor income.

Analytically characterizing the optimal consumption-savings choice (and thus the regression coefficients of income on consumption and wealth at different horizons N) is generally difficult in models with potentially binding borrowing constraints and/or preferences that deviate from quadratic utility. For important contributions dealing with borrowing constraints, see Holm (2018) and Carroll et al. (2021), for explicit solutions for CARA utility, see Caballero (1990) and Wang (2003). Note that for CARA utility, precautionary saving simply reduces consumption (and increases saving) by a constant in every period, leaving the regression coefficients at all horizons N completely unaffected.

4.3.1 Theoretical Predictions of a Simple Partial Insurance Framework

Rather than fully articulating an alternative model that can be solved either analytically or numerically and confronted with the empirical regression results, as in the important contribution by Fella et al. (2020), we now stipulate a simple extension to the PIH-implied consumption function that features partial consumption insurance against permanent shocks.²⁸ Blundell et al. (2008) find significant such partial insurance in the U.S. micro data (and Kaplan and Violante (2011) show that standard quantitative life cycle models with idiosyncratic income risk have difficulties reproducing the extent of that insurance). To this end, suppose that individual household consumption (and from the budget constraint, financial wealth) follow the rule

$$c_t = rw_t + (1 - \kappa)\eta_t + z_{t-1} + \frac{r\varepsilon_t}{1+r} \quad (11)$$

$$w_{t+1} = w_t + \kappa\eta_t + \frac{\varepsilon_t}{1+r} \quad (12)$$

where κ measures the degree of consumption insurance against permanent income shocks. The PIH is nested with $\kappa = 0$ (and so is complete insurance with $\kappa = 1$). Alternatively, one can interpret $1 - \kappa$ as the marginal propensity of current consumption out of a *permanent* income shock. We certainly do not argue that this is necessarily an optimal decision rule of a standard consumption-savings model with a precautionary motive (induced by either prudence or potentially binding borrowing constraints), but we do submit that it captures the most salient discrepancy between the PIH and the empirical record on the consumption response to income shocks.

In Appendix E, we show that the consumption and wealth regression coefficients in this augmented model are given by

$$\beta_c^N = \frac{[(1 - \kappa) + (N - 1)(r\kappa + 1)]Q + \frac{r}{1+r}(1 - M)}{NQ + 2} \quad (13)$$

$$\beta_w^N = \frac{(1 + r)\kappa NQ + (1 - M)}{(1 + r)[NQ + 2]} \quad (14)$$

These expressions imply (see again Appendix E) that while the consumption regression coefficients continue to be increasing in N , the wealth regression coefficients are declining in the horizon N only if insurance against permanent shocks is sufficiently imperfect, but are *increasing* in N if κ is sufficiently large. Concretely, the precise condition for β_w^N to be

²⁸ A similar approach is taken by Pedroni et al. (2023) in their study of the importance of advance information for consumption insurance, and by Ghosh and Theloudis (2023) in their study of partial consumption insurance against higher order income risk.

increasing (rather than decreasing) in N is

$$\kappa > \frac{(1 - M)}{2(1 + r)} \quad (15)$$

Evidently, this condition is always violated for the pure PIH with $\kappa = 0$, as must be the case since the PIH is a special case and we argue in the previous section that it implies declining (in N) wealth regression coefficients.

We now proceed in parallel to the previous section and interpret our empirical regression coefficients in light of this extended model. In particular, if we continue to use the $N = 2$ regression coefficients for wealth and consumption as targets for our 3-parameter model, but add to it the longer-run wealth response β_w^4 (since, as argued above, our simple consumption function implies the sharp condition (15) on the degree of permanent shock insurance κ) together with the degree of measurement error under which the wealth regression coefficients are increasing (or decreasing) with the horizon N . Equation (9) for $N = 2$ and equation (10) for $N = 2, 4$ constitute a system of three equations in three unknowns (Q, M, κ) that can be solved in closed form (see Appendix E). The implied parameter values and their empirical targets (replicated from Table 5) are displayed in Table 7.

4.3.2 Empirical Implementation

From Table 7, we observe that the values for the persistence of income shocks and measurement error deviate somewhat from those we obtain when restricting $\kappa = 0$. The estimate for the persistence of income shocks becomes larger at $Q = 0.74$, not surprisingly, since in order to match the same consumption response to income shocks with partial insurance against permanent shock requires a larger share of these income shocks to be permanent. Since more persistent income shocks call for a smaller wealth response, to rationalize the same wealth response over two periods requires, from the perspective of the model, a larger degree of measurement error now. The value for M rises from close to zero in the previous section to what we think is a more plausible value of $M = 0.5$

Table 7: Targets and Parameters

Target	Value	Param.	Value
β_c^2	0.347	Q	0.737
β_w^2	0.297	M	0.505
β_w^4	0.317	κ	0.364

Note: Empirical targets and model-implied parameters for the partial insurance framework specified in equations (11) and (12).

Most importantly, our estimate of the degree of insurance against permanent income shocks implied by the regression coefficients is $\kappa = 0.364$. That is, only 64% of a permanent income shock translates immediately into consumption whereas the rest (according to the model we stipulate) is temporarily insured and absorbed by wealth. Our estimate implies that inequality (15) is satisfied and thus in the partial insurance model (as in the data) the wealth regression coefficients continue to be mildly increasing with the horizon N , as in the original PIH.

Incidentally, our estimate of the extent of partial consumption insurance against permanent income shocks aligns very well with the consumption insurance coefficient for permanent shocks estimated by Blundell et al. (2008) of 36%, although it is important to note that the two numbers are not directly comparable since we study the consumption response in levels (by estimating regressions in first differences, motivated by the original PIH and the quadratic utility function that underlies it), whereas they, motivated by a first-order approximation of the stochastic Euler equation under CRRA utility, estimate a specification in log-differences and the insurance coefficients need to be interpreted as indicating what share of a one percent permanent income shock transmits (or does not transmit) into consumption growth. With our specification in levels we can interpret $1 - \kappa$ as the marginal propensity to consume out of a permanent income shock. Note that Kaplan and Violante (2011), in their study of a standard incomplete markets model, find a consumption insurance coefficient with respect to permanent shocks of 0.23, somewhat smaller than the value we infer from our Italian data and the simple consumption function stipulated in this section.

Table 8: Consumption Response

Statistic	Data	PIH	Par.In.	PIH*
β_c^2	0.347	0.347	0.347	0.424
β_c^4	0.331	0.515	0.542	0.597
β_c^6	0.412	0.615	0.648	0.689

Note: Consumption regression coefficients for different horizons N , data and models. The third column “Par.In.” is the partial insurance framework, and the last column is the PIH, but with the parameter estimates for Q, M implied by the partial insurance framework in Table 7.

The resulting partial insurance model-implied consumption- and wealth regressions over shorter and over longer horizons are contained in Tables 8 and 9, respectively (in the third columns in both tables), together with their empirical counterparts (the first column) and the numbers implied by the pure PIH in the second column (that is, we reproduce the information from Table 5 here for comparison). The last column labeled PIH* in both

tables draws out the implications of the standard PIH, but for the estimates of (Q, M) from this section.²⁹ Again note that the $N = 2$ consumption and wealth responses and now also the $N = 4$ wealth response were targeted in the calibration of (Q, M, κ) and thus it is no surprise that the partial insurance model can match the corresponding data moments perfectly.

Table 9 shows that the partial insurance model can replicate the longer run wealth response (for $N = 6$) almost perfectly. The fact that the model-implied wealth response is increasing with N is no surprise, given that the κ we infer from the data satisfies condition (15). However, the model matches the data along this dimension not only qualitatively, but also quantitatively well. The PIH, independent of whether one uses the estimates for (Q, M) from Section or 4.1.2 from this section, in contrast predicts a strongly monotonically declining pattern, qualitatively at odds with our estimates from the Italian data.

Table 9: Wealth Response

Statistic	Data	PIH	Par.In.	PIH*
β_w^2	0.297	0.297	0.297	0.143
β_w^4	0.317	0.221	0.317	0.101
β_w^6	0.316	0.175	0.328	0.077

Note: Wealth regression coefficients for different horizons N , data and models. The third column “Par.In.” is the partial insurance framework, and the last column is the PIH, but with the parameter estimates for Q, M implied by the partial insurance framework in Table 7.

Table 8 suggests that, qualitatively, all versions of the model considered in this paper are consistent with an increasing (with N) consumption response to income shocks, but that the gradient with respect of N is too steep relative to the data.³⁰ This in turn suggests that additional consumption insurance possibilities against income shocks accumulating over long time horizons exist in Italy other than what is implied by the simple partial insurance model introduced in this section (or the original PIH, for that matter).

Overall, we conclude from the results in this section that, qualitatively, a simple extension of the PIH consumption function that permits substantial insurance against permanent income shocks describes the consumption and wealth response to income shocks observed in the Italian SHIW data well, and that permitting the additional consumption insurance

²⁹ Equivalently, they are obtained by setting the value for partial insurance κ to zero but maintaining the estimates of the other two parameters (Q, M) from this section.

³⁰ The fact that the consumption response is stronger (the gradient steeper) in the partial insurance model than in the pure PIH (comparing the second and the third column of the table) stems from the fact that the estimated Q in this section is significantly larger than in the previous section. If one applies the same Q for the PIH (see the last column), this pattern is reversed.

against permanent shocks is crucial for successfully confronting the empirical predictions concerning the dynamic wealth responses.

It is important to note that these conclusions are derived from a sample of households without entrepreneurial and rental income, that is, from a (clearly non-representative) sample of households that we think should adhere best to the basic assumptions and predictions of the PIH and its simple extensions. In the next section, we will document that the consumption- and especially the wealth responses to income shocks of entrepreneurs and owner occupiers and landlords look quite different.

5 What Drives the Co-Movement between Income Shocks and Business and Real Estate Wealth?

5.1 The Role of Self-Employment


In our sample, a significant fraction of households (about 32%) report some business income. In this section, we establish that these households face larger income shocks, but also that they exhibit a different consumption- and wealth co-movement with these income shocks.

In Figure 5.1, we order households with respect to residual income changes, sort them into twenty equally sized bins and for each bin plot the fraction of households which report non-zero business income in at least one of the two years over which the income change is calculated. The figure clearly shows that households with business income experience, on average, larger absolute and relative (that is, logged) income changes.³¹ We then investigate whether and to what extent the consumption and wealth response of this group differs from the overall sample. To separate the effect of business wealth from that of real estate wealth, now we select all households which report no income from real estate in two consecutive years and we divide them into self-employed (reporting some business income in one of the years) and not self-employed (reporting zero business income in both years).

Line 1 of Table 10 reports the consumption (total and non-durable) and wealth co-movement with income shocks for this group of households. The last line of the table reports, for comparison, the co-movement for the group of households with neither business nor real estate wealth, which is the same group we used to produce the results in Table 4 above. We want to highlight that the consumption response for the households with business income is smaller than the one for the households without business income, suggesting that

³¹ Guiso et al. (2005) document that Italian firms provide substantial earnings insurance to their employees against firm-specific shocks. The stark difference between the earnings shocks for employees and self-employed in Figure 5.1 could therefore partly be due to the fact that employees are partially insured by their firms against idiosyncratic (to the firm or to the worker) productivity shocks.

Figure 7: Income changes and self-employment



Note: Households are sorted (by the size of the residual (log) labor income change) into 20 bins, and the y-axis reports the average share of self-employed within each bin.

the income process for the self-employed might be more volatile but less persistent (i.e., has a lower value for the parameter Q). A less persistent process would imply a stronger wealth response according to the PIH; however, the observed wealth response for the self-employed is more than 160 cents for a one Euro income shock, which is much larger than the one consistent with the Q implied by the consumption response.³² Moreover, the table shows that the large wealth response is driven primarily by the response in business wealth.

We conclude that households with business income face a more volatile and possibly less persistent income process but they also might face shocks to their business wealth that are correlated with their business income, for example, a persistent increase in the demand for the product of the business that raises the income of the business owners and, at the same time, the value of the business itself.

³² Using the non-durable consumption response of 18.9 cents and assuming zero interest rate, the estimated Q would be 0.23, and this would imply a wealth response of at most 0.46 (with zero measurement error) cents to the Euro.

Table 10: The Role of Self-Employment and Real Estate

	Δc	Δc_n	$\Delta \text{Real Wealth}$	$\Delta \text{FinWealth}$
1. Business income, no real estate Sample size: 703	27.2 (5.3)	18.9 (5.8)	129.6 (61.9)	32.3 (8.9)
2. Real estate, no business income Sample size: 10147	35.5 (3.1)	29.4 (2.8)	157.1 (26.6)	16.7 (10.0)
3. Real estate, no bus. income, non-adj. Sample size: 3359	38.8 (5.1)	29.8 (3.4)	63.3 (25.9)	-13.9 (29.7)
4. Business income and real estate Sample size: 5199	7.8 (3.4)	5.9 (1.3)	311.0 (70.8)	16.8 (11.4)
5. No business income, no real estate Sample size: 2612	36.0 (3.9)	34.7 (3.2)	NA	29.7 (4.4)

Note: Consumption and wealth regression coefficients for different samples of households, depending on whether the household has business income and/or real estate. Standard errors, clustered at the household level, are in parentheses. Real wealth in line 1 is business wealth, in lines 2 and 3 is real estate wealth, in line 4 is business plus real estate wealth.

5.2 Real Estate Wealth

In Italy, real estate is the predominant form of wealth held by private households. As shown in Table 1 above, in 2016 average net real estate wealth (that is the value of real estate minus mortgages) accounts for more than 75% of total net wealth. Moreover, real estate wealth is diffused. Around 69% of the households in our sample own their residence, and around 23% own more than one property. It is therefore not entirely surprising that adjustments in the real value of real estate wealth may play an important role in understanding the wealth and consumption response of households to an income shock.

As in the previous case, in order to separately identify the role of real estate wealth from that of business wealth we first look, in line 2 of Table 10, at the consumption and wealth response for households who own real estate but have no business income. For these households, the consumption response to income shocks is not significantly different from the response of similar households without real estate, suggesting that real estate per se does not change the extent of consumption insurance. However, real estate owners, like business owners, display a strong co-movement of wealth changes (over 170 cents to the Euro), primarily driven by changes in real estate wealth.

In order to better understand the sources of this large wealth response, in line 3 we restrict the sample of real estate owners that report the same exact portfolio of properties over the horizon of the income change. The line shows that even for non-adjusting households with positive real estate wealth there is a strong positive correlation between reported

income changes and reported real estate *price changes* of the continually owned properties. This correlation could possibly stem from a strong positive correlation of local housing and local labor markets.³³

We conclude this section by reporting, in line 4 of Table 10, the consumption and wealth responses for households who own both businesses and real estate. Not surprisingly, the wealth response to income shocks for these households is extremely large (over 300 cents to the Euro), as for these households the correlation of income shocks with *both* real estate and the value of business are at work. Perhaps more surprising is the very small (less than 10 cents to the Euro) consumption response to income shocks. One possibility is that for these households a positive income shock triggers an investment into their real estate and business (which are possibly connected), and these investments lower the corresponding consumption responses.

6 Conclusion

How do households respond to an income shock? In this paper we have answered this question using panel data from the Italian SHIW. We found that the responses are heterogeneous across groups. The consumption response of households who do not have a business is between 20 and 30 cents to the Euro, and it is substantially smaller (less than 10 cents to the Euro) for households who own a business. Wealth responses instead are around 30 cents to the Euro for households who do not own businesses or real estate, but substantially larger (exceeding 100 cents to the Euro) for households with businesses or real estate.

We have then argued that for the selected sample of households without real estate and business income a modified version of the standard permanent income hypothesis can account well for the wealth- and consumption response, both in the short (2 years) and long (6 years) run. Theory combined with data suggests that a significant share of income shocks are permanent in nature and that a substantial share of these shocks are insured in the short run. Interpreted from the perspective of our model that permits partial insurance even against permanent shocks, our estimates suggest a consumption response of only 64 cents for every Euro of a permanent shock; that is, 36 cents of the shock do *not* transmit into

³³ For the U.S., Davidoff (2006) documents a strong positive correlation between income growth and house price growth at the local level over five year horizons. He merges panel data on wages by region (MSA) and industry (2 digit SIC) from the BLS with regional (MSA) house price data from OFHEO and estimates an average (over MSA-industry pairs) correlation between house price and income growth of 0.29. The highest correlation (0.64) is obtained for households working in the amusement industry in the Orlando area. Davis and Ortalo-Magné (2011) find a strong positive correlation of 0.81 of mean (standardized) wage *levels* and rents in a year 2000 cross-section of MSA's.

current consumption. This accords well with the empirical estimates of Blundell et al. (2008) for the U.S. More broadly, this analysis demonstrates how observing both the consumption- and the wealth response to income shocks over shorter and longer time horizons (which requires panel data on all three variables) is informative about the consumption function (equivalently, the saving functions) characterizing household behavior.

The very strong wealth responses for housing and business owners suggest that shocks to the value of these assets are important in shaping household economic decisions, and that these shocks might be strongly correlated with labor income shocks faced by households, and thus constitute an important component of the resource risk faced by households. Future research should address in more detail the forces behind the large co-movement between income shocks and value of wealth, with the objective of developing a unified consumption-savings model that incorporates these idiosyncratic valuation shocks, and endogenizes the housing adjustment and business ownership decisions.

References

- Aaronson, D., S. Agarwal, and E. French (2012). The Spending and Debt Responses to Minimum Wage Hikes. *American Economic Review* 102(7), 3111–3139.
- Aiyagari, S. R. (1994). Uninsured Idiosyncratic Risk and Aggregate Saving. *Quarterly Journal of Economics* 109(3), 659–684.
- Altonji, J. and A. Siow (1987). Testing the Response of Consumption to Income Changes with (Noisy) Panel. *Quarterly Journal of Economics* 102(2), 293–328.
- Berger, D., V. Guerrieri, G. Lorenzoni, and J. Vavra (2017). House Prices and Consumer Spending. *Review of Economic Studies* 85(3), 1502–1542.
- Blundell, R., L. Pistaferri, and I. Preston (2008). Consumption Inequality and Partial Insurance. *American Economic Review* 98(5), 1887–1921.
- Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2016). Consumption Inequality and Family Labor Supply. *American Economic Review* 106(2), 387–435.
- Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2018). Children, Time Allocation, and Consumption Insurance. *Journal of Political Economy* 126(S1), S73–S115.
- Caballero, R. (1990). Consumption Puzzles and Precautionary Saving. *Journal of Monetary Economics* 25(1), 113–136.

- Campbell, J. and Z. Hercowitz (2009). Welfare Implications of the Transition to High Household Debt. *Journal of Monetary Economics* 56(1), 1–16.
- Carroll, C., M. Holm, and M. Kimball (2021). Liquidity Constraints and Precautionary Saving. *Journal of Economic Theory* 195, 1–21.
- Checchi, D., T. Jappelli, I. Marino, and A. Scognamiglio (2023). Inequality Trends in a Slow-Growing Economy: Italy 1990-2020. *CSEF Working Paper* 679.
- Cochrane, J. (1991). A Simple Test of Consumption Insurance. *Journal of Political Economy* 99(5), 957–976.
- Colarieti, R., P. Mei, and S. Stantcheva (2024). The how and why of household reactions to income shocks. *NBER Working Paper* 32191.
- Commault, J. (2022). Does Consumption Respond to Transitory Shocks? Reconciling Natural Experiments and Semistructural Methods. *American Economic Journal: Macroeconomics* 14(2), 96–122.
- Davidoff, T. (2006). Labor Income, Housing Prices, and Homeownership. *Journal of Urban Economics* 59(2), 209–235.
- Davis, M. and F. Ortalo-Magné (2011). Household Expenditures, Wages, Rents. *Review of Economic Dynamics* 14(2), 248–261.
- Deaton, A. (1992). *Understanding Consumption*. Oxford: Clarendon Press.
- Dynarski, S. and J. Gruber (1997). Can Families Smooth Variable Earnings? *Brookings Papers on Economic Activity* 1997(1), 229–284.
- Fagereng, A., M. B. Holm, B. Moll, and G. Natvik (2019). Saving Behavior Across the Wealth Distribution: The Importance of Capital Gains. *NBER Working Paper* 26588.
- Fella, G., S. Frache, and W. Koeniger (2020). Buffer-Stock Saving and Households' Response to Income Shock. *International Economic Review* 61(3), 1359–1382.
- Fernandez-Villaverde, J. and D. Krueger (2011). Consumption and Saving over the Life Cycle: How Important are Consumer Durables? *Macroeconomic Dynamics* 15(5), 725–770.
- Friedman, M. (1957). *A Theory of the Consumption Function*. Princeton University Press.

- Gervais, M. and P. Klein (2010). Measuring Consumption Smoothing in CEX Data. *Journal of Monetary Economics* 57(8), 988–999.
- Ghosh, A. and A. Theloudis (2023). Consumption Partial Insurance in the Presence of Tail Income Risk. *Working Paper*.
- Guiso, L., L. Pistaferri, and F. Schivardi (2005). Insurance within the Firm. *Journal of Political Economy* 113(5), 1054–1087.
- Guner, N., R. Kaygusuz, and G. Ventura (2012). Taxation and Household Labour Supply. *Review of Economic Studies* 79(3), 1113–1149.
- Guvenen, F. and A. Smith (2014). Inferring labor income risk and partial insurance from economic choices. *Econometrica* 82(6), 2085–2129.
- Hall, R. and F. Mishkin (1982). The sensitivity of consumption to transitory income: Estimates from panel data on households. *Econometrica* 50(2), 461–481.
- Heathcote, J., K. Storesletten, and G. Violante (2014). Consumption and labor supply with partial insurance: An analytical framework. *American Economic Review* 104(7), 2075–2126.
- Holm, M. (2018). Consumption with liquidity constraints: An analytical characterization. *Economics Letters* 167, 40–42.
- Holm, M. (2023). How Do Households Respond to Income Shocks?: Discussion. *available at: <https://drive.google.com/file/d/11ICYAsKaYMBFFCtFo1TCvVSg9LjHV-50/view>*.
- Holter, H., D. Krueger, and S. , Stepanchuk (2019). How Does Tax Progressivity and Household Heterogeneity Affect Laffer Curves? *Quantitative Economics* 10(4), 1317–1356.
- Hurst, E., A. Lusardi, A. Kennickel, and F. Torralba (2010). The importance of business owners in assessing the size of precautionary savings. *Review of Economics and Statistics* 92(1), 61–69.
- Jappelli, T., M. Padula, and L. Pistaferri (2008). A Direct Test of the Buffer-Stock Model of Saving. *Journal of the European Economic Association* 6(6), 1186–1210.
- Jappelli, T. and L. Pistaferri (2000a). The dynamics of household wealth accumulation in Italy. *Fiscal Studies* 21(2), 269–295.

- Jappelli, T. and L. Pistaferri (2000b). Using subjective income expectations to test for excess sensitivity of consumption to predicted income growth. *European Economic Review* 44(2), 337–358.
- Jappelli, T. and L. Pistaferri (2006). Intertemporal Choice and Consumption Mobility. *Journal of the European Economic Association* 4(1), 75–115.
- Jappelli, T. and L. Pistaferri (2010a). The Consumption Response to Income Changes. *Annual Review of Economics* 2, 479–506.
- Jappelli, T. and L. Pistaferri (2010b). Does Consumption Inequality Track Income Inequality in Italy? *Review of Economic Dynamics* 13(1), 133–153.
- Jappelli, T. and L. Pistaferri (2011). Financial Integration and Consumption Smoothing. *Economic Journal* 121(553), 678–706.
- Jappelli, T. and L. Pistaferri (2017). *The Economics of Consumption: Theory and Evidence*. Oxford University Press.
- Kaplan, G., B. Moll, and G. L. Violante (2018). Monetary Policy According to HANK. *American Economic Review* 108(3), 697–743.
- Kaplan, G. and G. L. Violante (2011). How Much Insurance Beyond Self-Insurance? *American Economic Journal: Macroeconomics* 2(4), 53–87.
- Kaplan, G. and G. L. Violante (2022). The marginal propensity to consume in heterogeneous agent models. *Annual Review of Economics* 14, 747–775.
- Kerr, S. P., W. Kerr, and T. Xu (2017). Personality traits of entrepreneurs: A review of the recent literature. *NBER Working Paper* 24097.
- Krueger, D., K. Mitman, and F. Perri (2016). Macroeconomics and Household Heterogeneity. Volume 2 of *Handbook of Macroeconomics*, pp. 843–921. Elsevier.
- Krueger, D. and F. Perri (2005). Understanding consumption smoothing: Evidence from the US Consumer Expenditure Survey. *Journal of the European Economic Association Papers and Proceedings* 3(2-3), 340–349.
- Krueger, D. and F. Perri (2006). Does income inequality lead to consumption inequality? Evidence and theory. *Review of Economic Studies* 73(1), 163–193.

- Luengo-Prado, M. and B. Sorensen (2008). What can explain excess smoothness and sensitivity of state-level consumption? *Review of Economics and Statistics* 90(1), 65–80.
- Mace, B. J. (1991). Full insurance in the presence of aggregate uncertainty. *Journal of Political Economy* 99(5), 928–956.
- Mazzocco, M. and S. Saini (2012). Testing efficient risk sharing with heterogeneous risk preferences. *American Economic Review* 102(1), 428–468.
- Padula, M. (2004). Consumer durables and the marginal propensity to consume out of permanent income shocks. *Research in Economics* 58(4), 319–341.
- Pedroni, M., S. Singh, and C. Stoltenberg (2023). Advance information and consumption insurance: Evidence and structural estimation. *Working Paper*.
- Schulhofer-Wohl, S. (2011). Heterogeneity and Tests of Risk Sharing. *Journal of Political Economy* 119(5), 925–958.
- Townsend, R. (1994). Risk and Insurance in Village India. *Econometrica* 62(3), 539–591.
- Wang, N. (2003). Caballero Meets Bewley: The Permanent-Income Hypothesis in General Equilibrium. *American Economic Review* 93(3), 927–936.

Appendix

A Variable Definitions and Panel Dimension

Non-durable consumption c_{nt} is defined as all household expenditures during a year, *minus* expenditures on transportation equipment (cars, bikes, etc.), valuables (such as art, jewelry, antiques), household equipment (such as furniture, rugs, TV's, cell phones, and other electronics), expenditure for home improvement, insurance premia, and contribution to pension funds. It includes rent paid by renters and imputed rent of homeowners on all properties that are not rented out. Imputed rent also appears as income from real assets on the right-hand side of the budget constraint. Expenditures on durables c_{dt} include expenditures for transportation equipment, valuables, and household equipment, all as defined above.

Labor income y_t is measured after taxes and includes fringe benefits received by employees and business income by entrepreneurs. Transfers T_t include both transfer payments from the government (such as unemployment benefits) as well as gifts, loans, and other transfers between private households.

Financial assets a_{t+1} add bank deposits, stock and bond holdings, and other direct holdings of financial assets (including assets held in private pension funds), net of outstanding debt. It does not include the value of entitlements to government pension payments. The net income from financial assets (interest payments, dividends, etc.) forms financial income. Finally, real assets e_{t+1} include the value of real estate property, the value of valuables (as defined above), and the net value of ownership in private businesses and partnerships. Income from real assets consists mainly of rent (both actual and imputed) received from owned real estate.

Table A1 documents the total sample size by year and the extent of the panel dimension of the SHIW as every column reports the total number of households interviewed in that year and then divides them by their entry year in the SHIW.

Table A1: SHIW sample size and panel dimension

	Year of interview												
	1991	1993	1995	1998	2000	2002	2004	2006	2008	2010	2012	2014	2016
Total sample	8188	8089	8135	7147	8001	8011	8012	7768	7977	7951	8151	8156	7420
By entry year:													
1991	8188	3470	2579	1713	1236	920	694	582	521	439	367	243	200
1993		4619	1066	583	399	270	199	157	141	124	106	78	59
1995			4490	373	245	177	117	101	84	75	62	46	37
1998				4478	1993	1224	845	636	538	450	380	267	207
2000					4128	1014	667	475	398	330	256	170	139
2002						4406	1082	672	525	416	340	221	161
2004							4408	1334	995	786	631	395	306
2006								3811	1143	856	648	414	298
2008									3632	1145	806	481	347
2010										3330	1015	579	385
2012											3540	1565	912
2014												3697	753
2016													3616

Note: Sample size of the SHIW, broken down by year of household entry. The second row gives the total number of observations for a given year, and the subsequent rows display, by year in which the household first entered the survey, the number of households from a given entry year still present in the sample at a (weakly) later year.

B Estimating Q from Income Data Alone

The income process in levels is given by

$$\begin{aligned}
 y_{it} &= z_{it} + \iota_{it} \\
 z_{it} &= z_{it-1} + \eta_{it}
 \end{aligned}$$

where ι_{it} is an i.i.d. random variable that includes temporary income shocks and measurement error. It is easy to show that

$$\begin{aligned}
 y_{it} - y_{it-2} &\equiv \Delta^2 y_{it} = z_{it} + \iota_{it} - z_{it-2} - \iota_{it-2} \\
 &= \iota_{it} + \eta_{it-1} + \eta_{it} - \iota_{it-2} \\
 y_{it+2} - y_{it} &\equiv F^2 y_{it} = \iota_{it+2} + \eta_{it+2} + \eta_{it-1} - \iota_{it}
 \end{aligned}$$

This implies that

$$\text{Var}(\Delta^2 y_{it}) = \text{Var}(F^2 y_{it}) = 2\text{Var}(\eta_{it}) + 2\text{Var}(\iota_{it})$$

and that

$$\text{cov}(\Delta^2 y_{it}, F^2 y_{it}) = -\text{Var}(\iota_{it})$$

B.1 Estimation Steps

- Select all households that are in the sample at $t-2, t, t+2$, and for these three periods have no income from real estate or business.
- Estimate $\text{Var}(\iota_{it})$ as $-\text{cov}(\Delta^2 y_{it}, F^2 y_{it}) \equiv A$
- Estimate $2(\text{Var}(\eta_{it}) + \text{Var}(\iota_{it}))$ as $(\text{Var}(\Delta^2 y_{it}) + \text{Var}(F^2 y_{it}))/2 \equiv B$. Note that in a large sample $\text{Var}(\Delta^2 y_{it})$ and $\text{Var}(F^2 y_{it})$ should be the same but in small sample they are not, so we average across those.
- Note that

$$\frac{B}{A} = 2(Q + 1)$$

so an estimate of $Q = \frac{\text{Var}(\eta_{it})}{\text{Var}(\iota_{it})}$ is given by

$$Q = \frac{1}{2} \frac{B}{A} - 1$$

B.2 Estimation Results

Using our sample yields an estimate of $Q = 0.60$.

C Predictable Income Changes

To the extent that our first-stage regression that conditions the data on observables such as age, education, etc., has failed to capture all predictable movements in income, the empirical estimates may partially reflect the consumption response to predictable income changes.³⁴ The PIH model of course implies that consumption should not respond to predictable changes in income at all. Denoting the predictable part of income by \bar{y}_t , the model now implies, for an income process,

$$\begin{aligned} \tilde{y}_t &= \bar{y}_t + z_t + \varepsilon_t + \gamma_t \\ z_t &= z_{t-1} + \eta_t \end{aligned}$$

³⁴ On the other hand, it is possible that some of the variation the first-stage regression picks up may have been predicted by the econometrician, but not by the household itself.

the following model solution

$$\begin{aligned}\Delta \tilde{y}_t &= \eta_t + \Delta \bar{y}_t + \Delta \varepsilon_t + \Delta \gamma_t \\ \Delta c_t &= \frac{r}{1+r} \varepsilon_t + \eta_t \\ \Delta w_{t+1} &= \frac{\varepsilon_t}{1+r} - \frac{1}{1+r} \sum_{s=1}^{\infty} \frac{\Delta \bar{y}_{t+s}}{(1+r)^{s-1}}.\end{aligned}$$

N -period changes are therefore given by

$$\begin{aligned}\Delta^N c_t &= \sum_{\tau=t-N+1}^t \left(\frac{r\varepsilon_\tau}{1+r} + \eta_\tau \right) \\ \Delta^N \tilde{y}_t &= \sum_{\tau=t-N+1}^t \eta_\tau + \Delta^N \bar{y}_t + \Delta^N \varepsilon_t + \Delta^N \gamma_t \\ \Delta^N w_{t+1} &= \sum_{\tau=t-N+1}^t \frac{\varepsilon_\tau}{1+r} - \frac{1}{1+r} \sum_{s=1}^{\infty} \frac{\Delta^N \bar{y}_{t+s}}{(1+r)^{s-1}}\end{aligned}$$

and the regression coefficients implied by the model now read as

$$\begin{aligned}\beta_c^N &= \frac{N\sigma_\eta^2 + r\sigma_\varepsilon^2/(1+r)}{N\sigma_\eta^2 + 2(\sigma_\varepsilon^2 + \sigma_\gamma^2) + Var(\Delta^N \bar{y}_t)} \\ \beta_w^N &= \frac{-\sum_{s=1}^{\infty} \frac{Cov(\Delta^N \bar{y}_t, \Delta^N \bar{y}_{t+s})}{(1+r)^s} + \sigma_\varepsilon^2/(1+r)}{N\sigma_\eta^2 + 2(\sigma_\varepsilon^2 + \sigma_\gamma^2) + Var(\Delta^N \bar{y}_t)}\end{aligned}$$

Thus the consumption response to income shocks goes down in the presence of predicted income changes, the extent to which is determined by how large the cross-sectional variance in the N -period change in the predictable component of income is, relative to the variance of the permanent and transitory income shocks. Note that the wealth response to income changes now also crucially depends on the covariance of current and future predicted income changes.

D Finite Lifetime

In the main paper we have assumed that households are infinitely lived, and have argued that for the question at hand this is an innocuous assumption for individuals far removed from their retirement age, while it is more problematic for individuals that are close to retirement. In this part of the appendix we provide the theoretical rationale for these

statements.

Now suppose that households have a finite lifetime horizon $T < \infty$, and for simplicity assume that T is known and constant across the population. In this case the optimal consumption rule under the PIH is given by

$$c_t = \frac{W_t}{\theta_{T-t}} \quad (16)$$

where

$$W_t = (1+r)w_t + E_t \sum_{s=0}^{T-t} \frac{y_{t+s}}{(1+r)^s} \quad (17)$$

are the expected lifetime resources remaining for the household, including initial wealth plus interest income from that wealth. The constant θ_{T-t} spreads out lifetime resources across the remaining lifetime and depends on the remaining lifetime (beyond the current period) $T-t$. It equals

$$\theta_{T-t} = \frac{1+r}{r} \left(1 - \frac{1}{(1+r)^{T-t+1}} \right) \quad (18)$$

if $r > 0$. Note that if $T \rightarrow \infty$, then $\theta_{T-t} = \frac{1+r}{r}$ and $c_t = \frac{r}{1+r} W_t$. If $r \rightarrow 0$ and $T < \infty$ then by L'Hopital's rule we have $\lim_{r \rightarrow 0} \theta_{T-t} = T-t+1$ (that is, remaining lifetime wealth, including expected human wealth, is simply divided by the number of remaining life years).

Given this consumption function, the realized changes in consumption and wealth in response to transitory and permanent shocks (ε_t, η_t) in a period $t \leq J$ during working life are given by

$$\Delta c_t = \left(\frac{1}{\theta_{T-t}} \right) \varepsilon_t + \left(\frac{\theta_{J-t}}{\theta_{T-t}} \right) \eta_t \quad (19)$$

$$\Delta w_{t+1} = \left(1 - \frac{1}{\theta_{T-t}} \right) \varepsilon_t + \left(1 - \frac{\theta_{J-t}}{\theta_{T-t}} \right) \eta_t \quad (20)$$

We note in particular that for $J = T$, then $\frac{\theta_{J-t}}{\theta_{T-t}} = 1$, and if furthermore $t = T$, then $\theta_{T-t} = 1$ as well. Finally, if $J < T$ but $t = J$, then θ_{J-t} , and in both cases there is no difference in the consumption (and asset) response to transitory and permanent income shocks. These are the consumption responses that we plot in Figure 6 of Section 4.2 in the main text.

The N -period consumption-, income-, and wealth changes consequently (for $t \leq J$)

become

$$\begin{aligned}
\Delta^N c_t &= \sum_{\tau=t-N+1}^t \left[\left(\frac{1}{\theta_{T-\tau}} \right) \varepsilon_\tau + \left(\frac{\theta_{J-\tau}}{\theta_{T-\tau}} \right) \eta_\tau \right] \\
\Delta^N w_{t+1} &= \sum_{\tau=t-N+1}^t \left[\left(1 - \frac{1}{\theta_{T-\tau}} \right) \varepsilon_\tau + \left(1 - \frac{\theta_{J-\tau}}{\theta_{T-\tau}} \right) \eta_\tau \right] \\
\Delta^N \tilde{y}_t &= \sum_{\tau=t-N+1}^t \eta_\tau + \Delta^N \varepsilon_t + \Delta^N \gamma_t
\end{aligned} \tag{21}$$

Finally, the bi-variate regression coefficients of N -period consumption and wealth changes on N -period income change are given as

$$\begin{aligned}
\beta_c^N &= \frac{\sum_{\tau=t-N+1}^t \left[\left(\frac{\theta_{J-\tau}}{\theta_{T-\tau}} \right) \right] \sigma_\eta^2 + \sigma_\varepsilon^2 / \theta_{T-t}}{N\sigma_\eta^2 + 2(\sigma_\varepsilon^2 + \sigma_\gamma^2)} \\
\beta_w^N &= \frac{\sum_{\tau=t-N+1}^t \left[\left(1 - \frac{\theta_{J-\tau}}{\theta_{T-\tau}} \right) \right] \sigma_\eta^2 + \left(1 - \frac{1}{\theta_{T-t}} \right) \sigma_\varepsilon^2}{N\sigma_\eta^2 + 2(\sigma_\varepsilon^2 + \sigma_\gamma^2)}.
\end{aligned}$$

D.1 No Retirement, $J = T$

In this case, we have $\frac{\theta_{J-t}}{\theta_{T-t}} = 1$, and the results are essentially unchanged from the infinite horizon case as long as $T - t$ is sufficiently large (i.e., households are sufficiently young) and thus $\theta_{T-t} \approx \frac{1+r}{r}$. See the left panel of Figure 6 in the main text.

D.2 With Retirement, $J < T$

The main difference to the previous case is that $\frac{\theta_{J-t}}{\theta_{T-t}} < 1$, and thus the consumption response to “permanent” shocks η_τ (which in the case of retirement are not really permanent) is smaller, and wealth response is larger. However, as long as households are young and thus t is small relative to J and T , the ratio $\frac{\theta_{J-t}}{\theta_{T-t}} \approx 1$ (see the left panel of Figure 6 in the main text) and the sum in the expression for the regression coefficients $\sum_{\tau=t-N+1}^t \left[\left(\frac{\theta_{J-\tau}}{\theta_{T-\tau}} \right) \right]$ will be dominated by terms close to 1, and in turn dominates the term in front of the transitory shock variance.

E Partial Insurance against Permanent Income Shocks: Details

First, recall that true income is given by the process

$$y_t = z_t + \varepsilon_t \quad (22)$$

$$z_t = z_{t-1} + \eta_t \quad (23)$$

with initial value $z_{-1} = \bar{y} \geq 0$. Measured income in the data is given by

$$\tilde{y}_t = y_t + \gamma_t = z_t + \varepsilon_t + \gamma_t \quad (24)$$

The budget constraint reads as

$$c_t + w_{t+1} = y_t + (1+r)w_t \quad (25)$$

with initial condition $w_0 \geq 0$. In the main text, we postulate a consumption rule of the form

$$c_t = rw_t + (1-\kappa)\eta_t + z_{t-1} + \frac{r\varepsilon_t}{1+r} \quad (26)$$

Together with the budget constraint, it implies the following realized changes in consumption and financial wealth of the form

$$\begin{aligned} \Delta w_{t+1} &= \kappa\eta_t + \frac{\varepsilon_t}{1+r} \\ \Delta c_t &= c_t - c_{t-1} \\ &= rw_t + (1-\kappa)\eta_t + z_{t-1} + \frac{r\varepsilon_t}{1+r} \\ &\quad - z_{t-1} - \varepsilon_{t-1} - (1+r)w_{t-1} + w_t \\ &= (1+r)\Delta w_t + (1-\kappa)\eta_t + \frac{r\varepsilon_t}{1+r} - \varepsilon_{t-1} \\ &= (1+r) \left(\kappa\eta_{t-1} + \frac{\varepsilon_{t-1}}{1+r} \right) + (1-\kappa)\eta_t + \frac{r\varepsilon_t}{1+r} - \varepsilon_{t-1} \\ &= (1+r)\kappa\eta_{t-1} + (1-\kappa)\eta_t + \frac{r\varepsilon_t}{1+r} \end{aligned}$$

E.1 N -Period Changes

Straightforward calculations imply that

$$\begin{aligned}
\Delta^N w_{t+1} &= \sum_{\tau=0}^{N-1} \left(\kappa \eta_{t-\tau} + \frac{\varepsilon_{t-\tau}}{1+r} \right) \\
\Delta^N c_t &= (1+r)\kappa \eta_{t-1} + (1-\kappa)\eta_t + \frac{r\varepsilon_t}{1+r} \\
&\quad + (1+r)\kappa \eta_{t-2} + (1-\kappa)\eta_{t-1} + \frac{r\varepsilon_{t-1}}{1+r} \\
&\quad + (1+r)\kappa \eta_{t-3} + (1-\kappa)\eta_{t-2} + \frac{r\varepsilon_{t-2}}{1+r} \dots \\
&\quad + (1+r)\kappa \eta_{t-N} + (1-\kappa)\eta_{t-N-1} + \frac{r\varepsilon_{t-N-1}}{1+r} \\
&= (1-\kappa)\eta_t + (r\kappa+1) \sum_{\tau=1}^{N-1} \eta_{t-\tau} + (1+r)\kappa \eta_{t-N} + \left(\frac{r}{1+r} \right) \sum_{\tau=0}^{N-1} (\varepsilon_{t-\tau}) \\
\Delta \tilde{y}_t^N &= \sum_{\tau=0}^{N-1} \eta_{t-\tau} + (\varepsilon_t + \gamma_t) - (\varepsilon_{t-N} + \gamma_{t-N})
\end{aligned}$$

E.2 Regression Coefficients

From the above expressions we can calculate the regression coefficients as

$$\begin{aligned}
\beta_c^N &= \frac{Cov(\Delta^N c_t, \Delta^N \tilde{y}_t)}{Var(\Delta^N y_t)} = \frac{[(1-\kappa) + (N-1)(r\kappa+1)]\sigma_\eta^2 + \frac{r}{1+r}\sigma_\varepsilon^2}{N\sigma_\eta^2 + 2(\sigma_\varepsilon^2 + \sigma_\gamma^2)} \\
\beta_w^N &= \frac{Cov(\Delta^N w_t, \Delta^N \tilde{y}_t)}{Var(\Delta^N y_t)} = \frac{(1+r)\kappa N\sigma_\eta^2 + \sigma_\varepsilon^2}{(1+r)[N\sigma_\eta^2 + 2(\sigma_\varepsilon^2 + \sigma_\gamma^2)]}
\end{aligned}$$

With $Q = \frac{\sigma_\eta^2}{\sigma_\varepsilon^2 + \sigma_\gamma^2}$ and $M = \frac{\sigma_\gamma^2}{\sigma_\varepsilon^2 + \sigma_\gamma^2}$, defined as in the main text, we can rewrite these expressions as

$$\beta_c^N = \frac{[(1-\kappa) + (N-1)(r\kappa+1)]Q + \frac{r}{1+r}(1-M)}{NQ + 2} \quad (27)$$

$$\beta_w^N = \frac{\kappa NQ + \frac{(1-M)}{1+r}}{[NQ + 2]} \quad (28)$$

E.3 Changes over the Horizon

Now we want to derive predictions about how the regression coefficients of consumption and wealth on income change with the time horizon N . To simplify the exposition, let us treat the time horizon as a continuous variable and take the derivative of (27) and (28)

with respect to N . This yields for the wealth regression coefficient

$$\begin{aligned}\frac{\partial \beta_w^N}{\partial N} &= (1+r)Q \frac{(1+r)\kappa [NQ+2] - [(1+r)\kappa NQ + (1-M)]}{((1+r)[NQ+2])^2} \\ &= \frac{Q}{(1+r)(NQ+2)^2} [2(1+r)\kappa - (1-M)]\end{aligned}$$

Assuming $Q > 0$ (that is, in the presence of at least some permanent income shocks), this expression is positive if and only if

$$\kappa > \frac{(1-M)}{2(1+r)} \quad (29)$$

that is, if and only if the degree of insurance against permanent shocks is sufficiently strong. The standard PIH is nested, with $\kappa = 0$, and thus β_w^N is unambiguously decreasing in the time horizon N . A flat wealth profile requires $\kappa = \frac{1-M}{2(1+r)}$.

For the consumption regression coefficient we find that

$$\begin{aligned}\frac{\partial \beta_c^N}{\partial N} &= \frac{(r\kappa+1)Q(NQ+2) - Q \left([(1-\kappa) + (N-1)(r\kappa+1)]Q + \frac{r}{1+r}(1-M) \right)}{(NQ+2)^2} \\ &= \frac{(1+r\kappa)(NQ^2+2Q) - [(1-\kappa) + (1+r\kappa)(N-1)]Q^2 - \frac{r(1-M)Q}{1+r}}{(NQ+2)^2} \\ &= \frac{(1+r\kappa)NQ^2 + (1+r\kappa)2Q - (1-\kappa)Q^2 - (1+r\kappa)(N-1)Q^2 - \frac{r(1-M)Q}{1+r}}{(NQ+2)^2} \\ &= Q \frac{(1+r)\kappa Q + (1+r\kappa)2 - \frac{r(1-M)}{1+r}}{(NQ+2)^2} > 0\end{aligned}$$

and thus the consumption regression coefficient is unambiguously increasing in N . The reason the wealth response is potentially increasing in the horizon N is as follows. In each period, a fraction κ of the current permanent shock η_t transmits into wealth. In the absence of transitory shocks, since both the covariance between income and wealth and the variance of income increase linearly with N , the regression coefficient would equal κ independent of N . But with transitory shocks the variance of income and the covariance between income and wealth grow with N at different rates, and the covariance grows with N the faster the larger is κ . Thus, if κ is sufficiently large, the wealth regression coefficients are increasing in N .

E.4 Determining the Parameters

E.4.1 Standard PIH

The standard PIH is a special case with $\kappa = 0$. We can use either (β_c^2, β_w^2) or (β_c^2, β_c^4) to estimate the two parameters, but in both cases the regression coefficient estimates have to obey bounds in order to deliver plausible values for (Q, M) . When using (β_c^2, β_w^2) , the estimates are

$$Q = \frac{\beta_c^2 - r\beta_w^2}{1 - \beta_c^2 + r\beta_w^2}$$

$$M = 1 - \frac{2(1+r)\beta_w^2}{1 - \beta_c^2 + r\beta_w^2}.$$

To obtain a plausible parameter value for $M \geq 0$ the restriction

$$(2+r)\beta_w^2 + \beta_c^2 \leq 1$$

which is satisfied unless both β_c^2 and β_w^2 are estimated to be too large. If we use consumption coefficients for $N = 2, 4$, we obtain

$$\beta_c^2 = \frac{2Q + \frac{r}{1+r}(1-M)}{2Q+2}$$

$$\beta_c^4 = \frac{4Q + \frac{r}{1+r}(1-M)}{4Q+2}$$

and thus the estimates for (Q, M) are given by

$$Q = \frac{2(\beta_c^2 - \beta_c^4)}{4\beta_c^4 - 2 - 2\beta_c^2} = \frac{2(\beta_c^4 - \beta_c^2)}{2 + 2\beta_c^2 - 4\beta_c^4}$$

$$M = 1 - \frac{1+r}{r} (\beta_c^2(2Q+2) - 2Q)$$

and again restrictions such as $\beta_c^4 > \beta_c^2$ and on the size of β_c^2 are required to be satisfied such that $Q, M \in [0, 1]$.

E.4.2 Partial Insurance Model

Now consider the extended model, for which

$$\begin{aligned}\beta_c^N &= \frac{[(1 - \kappa) + (N - 1)(r\kappa + 1)]Q + \frac{r}{1+r}(1 - M)}{NQ + 2} \\ \beta_w^N &= \frac{(1 + r)\kappa NQ + (1 - M)}{(1 + r)[NQ + 2]}\end{aligned}$$

If we assume, as in the main text, that $r = 0$ (the results for positive but small r are quantitatively very similar) then

$$\begin{aligned}\beta_c^N &= \frac{[(1 - \kappa) + (N - 1)]Q}{NQ + 2} \\ \beta_w^N &= \frac{\kappa NQ + (1 - M)}{[NQ + 2]}\end{aligned}$$

We continue to choose the $N = 2$ regression coefficients since doing so facilitates the comparison with the original PIH. Thus, we use

$$\begin{aligned}\beta_c^2 &= \frac{Q[2 - \kappa]}{2Q + 2} \\ \beta_w^2 &= \frac{2\kappa Q + (1 - M)}{2Q + 2} \\ \beta_w^N &= \frac{N\kappa Q + (1 - M)}{NQ + 2}\end{aligned}$$

for some $N \in \{4, 6, \dots\}$. For a given choice of N , these are three equations in the three unknowns (Q, M, κ) which can be solved numerically. Alternatively, the three equations become linear equations with a change of variables to $Q, M, \kappa Q$, which delivers the solution

$$\begin{aligned}Q &= \frac{2[\beta_w^N + (N - 2)\beta_c^2 - \beta_w^2]}{2\beta_w^2 + 2(N - 2)(1 - \beta_c^2) - N\beta_w^N} \\ \kappa &= 2 - \left(2 + \frac{2}{Q}\right)\beta_c^2 \\ M &= 1 + 2\kappa Q - (2Q + 2)\beta_w^2\end{aligned}$$

F Housing in the Standard Incomplete Markets Model

We now introduce housing explicitly into the standard incomplete markets model. We first model the housing choice of households without any frictions in the adjustment of

real estate position and no explicit borrowing constraints.³⁵ Also, households have access to a competitive rental market where housing services s_t can be rented for a rental price R_t per unit of house. Households buy real estate h_{t+1} at price per unit of p_t , as well as nondurable consumption c_{nt} and financial assets a_{t+1} . Houses depreciate at rate δ . The household decision problem is then given by

$$\begin{aligned} & \max_{\{c_{nt}, s_t, a_{t+1}, h_{t+1}\}} E_0 \sum_t \beta^t v(c_{nt}, s_t) \\ c_{nt} + a_{t+1} + R_t s_t + p_t h_{t+1} &= y_t + (1 + r_t) a_t + p_t (1 - \delta) h_t + R_t h_t \end{aligned} \quad (30)$$

where $v(c_{nt}, s_t)$ gives the period utility from consuming nondurables c_{nt} and housing services s_t .

F.1 Analysis

It is straightforward to show that this household problem can be solved in three stages. In the first stage, the intratemporal consumption allocation problem between non-durables and housing services is solved

$$\begin{aligned} & \max_{c_{nt}, s_t} v(c_{nt}, s_t) \\ c_{nt} + R_t s_t &= c_t \end{aligned}$$

where c_t is consumption expenditures on housing services. The solution characterized by the two equations

$$\begin{aligned} \frac{v_s(c_{nt}, s_t)}{v_{c_n}(c_{nt}, s_t)} &= R_t \\ c_{nt} + R_t s_t &= c_t \end{aligned}$$

Define the indirect utility function resulting from this maximization problem as

$$u(c_t; R_t) = v(c_{nt}(c_t, R_t), s_t(c_t, R_t))$$

This is the period utility function used in the main text.

In the second stage, the household decides how to split her savings between financial and real assets. Without any frictions in the real estate market (or the financial asset market, for that matter) a simple no-arbitrage argument implies that the rental price and the price

³⁵ Of course an appropriate no-Ponzi condition has to be imposed to make the household problem have a solution.

of real estate have to satisfy the condition.

$$R_{t+1} = p_t \left[(1 + r_{t+1}) - \frac{p_{t+1}(1 - \delta)}{p_t} \right]$$

Under this condition, one can consolidate both assets into one

$$w_{t+1} = a_{t+1} + p_t h_{t+1}.$$

Exploiting the outcome of the first two steps, the intertemporal household problem then reads as

$$\begin{aligned} \max_{\{c_t, w_{t+1}\}} E_0 \sum_t \beta^t u(c_t; R_t) \\ c_t + w_{t+1} = y_t + (1 + r_t)w_t \end{aligned}$$

where consumption expenditures and wealth are measured as

$$\begin{aligned} c_t &= c_{nt} + R_t s_t \\ w_{t+1} &= a_{t+1} + p_t h_{t+1} \\ &= a_{t+1} + e_{t+1}. \end{aligned}$$

As long as c_t and w_t are measured empirically consistent with the theory, the analysis can proceed as in the main text, without explicit consideration of the households' housing choice.

F.2 Adding Financing Constraints

Suppose the household can only finance a fraction $1 - \gamma$ of the value of real estate purchased in the current period,

$$a_{t+1} \geq -(1 - \gamma)p_t h_{t+1}.$$

The effect of such a constraint on the dynamics of the stock of real estate was studied, among others, by Fernandez-Villaverde and Krueger (2011), Campbell and Hercowitz (2009), and Aaronson et al. (2012). The presence of such a constraint may significantly alter the response of housing wealth to a change in income. Suppose that households find it optimal to be at the constraint in period t , then $a_{t+1} = -(1 - \gamma)p_t h_{t+1}$. Substituting this into the budget

constraint (30) yields

$$c_t + \gamma p_t h_{t+1} = y_t + (1 + r_t)a_t + p_t(1 - \delta)h_t + R_t h_t.$$

Therefore if households are constrained in periods $t - 1$ and t we have

$$\frac{\Delta c_t}{\Delta y_t} + \frac{\gamma \Delta p_t h_{t+1}}{\Delta y_t} = 1.$$

It is straightforward to observe $\frac{\Delta p_t h_{t+1}}{\Delta y_t} > 1$, that is expenditures on non-durables and net new housing can exceed the income change since households can leverage home purchases. But also note that this implies that

$$\frac{\Delta a_{t+1}}{\Delta y_t} = -(1 - \gamma) \frac{\Delta p_t h_{t+1}}{\Delta y_t}$$

and thus one would expect large adjustments in the value of mortgages (or other financial debt) too. This is not what the empirical analysis in Section 5 of the main text reveals.

F.3 Adding Prohibitive Transaction Costs

Now imagine a household with the current size of housing stock h that will never move, perhaps because of prohibitively high transaction costs. Also assume that the household lives in her own home, and does not rent out part of the home. Furthermore, assume that the depreciation rate on houses is $\delta = 0$. Let the service flow from the owner-occupied housing h be given by $s = \phi(h)$, where ϕ is an arbitrary function. Then the household problem (absent financing constraints) reads as

$$\begin{aligned} \max_{\{c_{nt}, a_{t+1}\}} E_0 \sum_t \beta^t v(c_{nt}, s) \\ c_{nt} + a_{t+1} + p_t(h_{t+1} - h_t) = y_t + (1 + r_t)a_t \end{aligned}$$

and since, by assumption, $h_{t+1} = h_t$, the budget constraint reads as

$$c_{nt} + a_{t+1} = y_t + (1 + r_t)a_t.$$

Evidently, the household does not care at all about changes in house prices. As long as the utility function is additively separable in c_n and s or satisfies

$$v(c_{nt}, s) = g(s)u(c_{nt}) = g(\phi(h))u(c_{nt}),$$

since h is constant by assumption, the presence of housing services simply represents an affine transformation of the period utility function that leaves non-durable consumption choice behavior unaffected. Thus, under these assumption we can proceed with our analysis of the PIH or the more elaborated precautionary savings model as if housing wealth and services are not present in the model.

With a nonzero depreciation rate $\delta > 0$, house price changes affect measured disposable labor income $\tilde{y}_t = y_t - p_t\delta h$, but the model-implied map between the adjusted income measure and non-durable consumption and financial wealth remains unaffected.