

Savoia, Ettore

Working Paper

The effects of labor income risk heterogeneity on the marginal propensity to consume

ECB Working Paper, No. 2866

Provided in Cooperation with:

European Central Bank (ECB)

Suggested Citation: Savoia, Ettore (2023) : The effects of labor income risk heterogeneity on the marginal propensity to consume, ECB Working Paper, No. 2866, ISBN 978-92-899-6243-8, European Central Bank (ECB), Frankfurt a. M., <https://doi.org/10.2866/07882>

This Version is available at:

<https://hdl.handle.net/10419/297306>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



EUROPEAN CENTRAL BANK

EUROSYSTEM

Working Paper Series

Ettore Savoia The effects of labor income risk
heterogeneity on the marginal
propensity to consume

No 2866

Abstract

Using detailed micro-data, this paper documents that households with lower income risk (and higher income levels) exhibit a higher Marginal Propensity to Consume (MPC) in response to transitory income shocks, all else being equal. This finding is particularly significant among unconstrained households and supported by models with precautionary saving only if designed to account for the empirically observed negative correlation between income levels and income risk. This interaction generates saving dynamics such that the stationary distribution of wealth among households facing different risk levels is not polarized. Therefore, it is possible to compare their respective MPCs within wealth and identify the reduction in MPC due to labor income risk. Otherwise, the effects of income risk are masked by wealth effects. In neither case, the MPC depends on (permanent, persistent, or current) income levels, whose direct effect on the MPC is always ambiguous. Finally, simulations of targeted fiscal rebates for specific labor categories reveal that governments cannot simultaneously stimulate aggregate demand and mitigate income risk.

JEL: D12, D52, D81, E21, J31

Key words: Marginal Propensity to Consume, Income Risk, Precautionary Saving

Non-technical Summary

The marginal propensity to consume (MPC) tells us how much a household consumes, within a given period, out of some unexpected change in income. This parameter is essential because understanding its strength allows policymakers to quantify the consumption reaction to shocks and policy changes that create a shift in a household's disposable income. Therefore, the MPC crucially shapes the effectiveness of fiscal and monetary policy.

Usually, there are different types of income shocks. Some can be transitory; others can last for longer or permanently. While some surprises are well-known in advance (e.g., work retirement, increase in education expenditure for children), sometimes, these shocks are unanticipated by households. Therefore, their consumption and saving plans are not adjusted in advance. This surprise allows economists to identify the household's consumption reaction due to a particular type of shock—the MPC. Finally, although unexpected and transitory, some shocks can be large, and others can be smaller and regardless of the size, shocks are either positive or negative.

According to the type of shock considered, we have different concepts of MPC. In this paper, I focus on shocks that are small, unexpected, and transitory. Further, I do not make any explicit distinction about their sign. Intuitive examples of these shocks are temporary tax or interest rate changes.

This paper provides new evidence to further our understanding of the effects of labor income risk heterogeneity on the MPC. Therefore starts by analyzing a simple question: Do households facing different income risk levels have different MPCs?

This question is important because the MPC shapes the effectiveness of macro policies, as argued above, and because people are heterogeneously exposed to income risk. For example, someone is employed under a permanent contract and thus has a stable income path. Other people instead have less stable jobs and face much more uncertainty concerning their future earnings. Therefore, it is important to understand how these structural sources of heterogeneity affect the household's consumption response out of transitory shocks.

However, the answer to this simple question is complex. The MPC depends on the level of disposable wealth. Affluent households can better afford small, unexpected, and transitory shocks than households with a low liquid wealth buffer. Accordingly, an affluent household's consumption change will be lower than that of a less wealthy household. Considering this problem from a risk perspective, a household facing more uncertainty can accumulate more savings (liquid wealth) than a household facing lower uncertainty.

Therefore, we do not know whether his MPC is lower because of a wealth effect or because of a risk effect.

To solve this issue and provide a clear answer, throughout this paper, I propose a quite intuitive strategy: It consists of comparing households that while holding the same level of liquid wealth (or same portfolio composition), cope against different levels of income risk. This strategy, adequately implemented using panel data from the Bank of Italy Survey of Household Income and Wealth (SHIW), allows for a shutdown of the wealth effects on the MPC and isolates the impact of income risk on the MPC.

As we will see in the paper, a critical element that allows households facing different levels of income risk to have almost the same amount of liquid wealth is given by the empirical fact that income risk negatively correlates with income levels. This correlation leads to an almost flat relationship between liquid assets and income risk and allows for comparing MPCs within the wealth distribution.

The main finding of this paper is that there exists an important degree of MPC heterogeneity within the wealth distribution, particularly among unconstrained, where given the same level of liquid resources, households facing a lower level of income risk (*safe*) display larger MPCs, 0.48 on average, compared to households facing a higher level of income risk (*risky*), who exhibit an average MPC of 0.28.

Intuitively, from the perspective of an unconstrained household whose income stream is predictable, there are few incentives to adjust saving after a small windfall. In contrast, when the uncertainty band is larger (*risky*), and households can adjust their consumption path, the income risk effect dominates. Therefore when households are far away from the borrowing constraint, the higher the risk, the smaller the part of the transitory shock transmitted to consumption. Similar conclusions cannot be reached for constrained households. Having more or less income risk might not be of first-order importance simply because the household cannot adjust consumption and saving decisions according to the risk it faces.

The overall evidence indicates that the Italian economy has a considerable fraction (25-38%) of unconstrained households facing low uncertainty and exhibiting higher biannual MPCs (out of transitory shocks) than households facing higher income risk. This fact has been formally rationalized through the lens of standard models with precautionary saving.

Finally, the evidence found in this paper highlights challenges for the design of targeted stimulus checks aimed at restoring particular labor categories facing substantial labor uncertainty. In these cases, a government faces a trade-off between mitigating income risk and maximizing the aggregate demand response of the fiscal expansion.

Introduction

The marginal propensity to consume (MPC) is a crucial parameter in heterogeneous agent macroeconomics. In this context, a standard definition of MPC is the change in non-durable consumption, within a given period, out of a transitory and unanticipated change in income, (Kaplan and Violante 2021). Following this definition, this paper provides empirical and theoretical evidence of MPC heterogeneity depending on households' exposure to labor income uncertainty.

While it is well known that labor income uncertainty raises the MPC at all levels of wealth compared to the case in which there is no uncertainty, i.e., certainty equivalence, (Zeldes 1989,a; Kimball 1990b), still the effects of *heterogeneity* in labor income risk exposure on the MPC are not clearly quantified and fully understood. This knowledge gap is of utmost importance given that households are heterogeneously exposed to labor income fluctuations and given that the MPC crucially shapes the effectiveness of the fiscal and monetary policy, (Kaplan, Moll, and Violante 2018). To fill this gap, I analyze a simple question: Do households facing different income risk levels have different MPCs? It turns out that the answer to this question is neither theoretically nor empirically straightforward.

In theory, whenever the utility function exhibits prudence, $U'''(\cdot) > 0$, income risk affects the MPC mainly through precautionary saving. In any heterogeneous agent models with precautionary saving, there are two offsetting forces: On the one hand, higher uncertainty in the current period increases savings and therefore decreases current consumption. This force makes the consumption function more concave¹ and raises the MPC at all levels of wealth. On the other hand, a higher level of saving implies that fewer households are close to the borrowing constraint in the stationary distribution, lowering the MPC². Therefore, when comparing the MPC between two uncertain consumption functions, a model with idiosyncratic income risk and a precautionary savings motive cannot accurately provide a prediction, (Zeldes 1989,a).

From an empirical point of view, there are two challenges: First, existing measures of income risk tend to be problematic either because they fail in capturing the income risk that cannot be avoided by households (Fagereng, Guiso, and Pistaferri 2016) or because they produce conditional moments that are related to income risk but fail to target risk

¹Note that the concavity of the consumption function is achieved not only when there is prudence, $U'''(\cdot) > 0$. The consumption function becomes concave even without a precautionary motive if there are occasionally binding borrowing constraints, (Carroll, Holm, and Kimball 2021). Further, the consumption function is concave when there is prudence even if there are no liquidity constraints, (Kaplan and Violante 2021).

²(Kaplan and Violante 2021) provides a clear exposition of these offsetting forces.

directly (Arellano et al. 2022). The second problem is related to the role of liquid assets. On the one hand, the vast empirical literature on the role of income risk on individual financial decisions and portfolio choices ³ documents the existence of a positive precautionary saving in the presence of higher uncertainty. On the other hand, there is significant evidence of the fact that the MPC is decreasing in liquid assets, (Jappelli and Pistaferri 2010b, 2014, 2020; Ganong et al. 2020; Fagereng, Holm, and Natvik 2018). This implies that any empirical test not able to hold constant the level of liquid assets cannot provide clear evidence of the effect of Labor income risk on the MPC.

To solve this issue and provide a clear answer, throughout this paper, I propose a quite intuitive strategy: It consists of comparing households that, while holding the same level of liquid assets, cope against different levels of income risk. Empirically, I use panel data from the Bank of Italy Survey of Household Income and Wealth, SHIW - hereafter. I proceed in a few steps: First, to correctly identify income risk, I follow (Arellano et al. 2022) and construct a Coefficient of Variation measure of *total* unpredictable income risk, based on expected income flows, at the individual level, *CV* -hereafter. I then split the sample between *risky* households, those with labor income risk above the median of the distribution, and *safe* households (below median risk). Second, by taking the household's portfolio composition as an observable variable and using the (Blundell, Pistaferri, and Preston 2008) estimator (*BPP* hereafter), I recover the MPC out of a transitory shock for households that have the same amount of wealth or, the same portfolio composition but lie at different points of the income risk distribution. Finally, I test the difference across groups.

This within-wealth comparison approach requires that the wealth distribution among the two masses of households is not polarized. That is, there must be sufficient masses of safe and risky households at each point of the wealth distribution. As illustrated later, it turns out that this is true in the data and can be rationalized by a well-known fact: the observed negative correlation between income level and income risk⁴. Interestingly, in the SHIW, the CV negatively correlates with actual, lagged, and permanent income. I show that this fact is empirically and theoretically key to generating homogeneity in the liquid assets holdings between the two groups of households. Once the two groups hold

³See: (Pratt and Zeckhauser 1987; Kimball 1990a; Guiso, Jappelli, and Terlizzese 1992, 1996; Bertaut and Haliassos 1997; Koo 1998; Heaton and Lucas 2000; Deaton 1991; Fagereng, Guiso, and Pistaferri 2016) among others.

⁴(Arellano et al. 2022), using administrative data from Spain shows that current and lagged income is negatively correlated with the CV. This relationship seems to hold even for different measures of income risk. Using administrative data from the US, (Güvenen et al. 2021) find that the standard deviation of income growth is decreasing in labor income once accounting for unemployment shocks.

almost the same amount of liquid assets, it is possible to identify within-wealth variations in MPC due to labor income risk. This paper makes four contributions:

- (a) Empirically, I uncover an important degree of MPC heterogeneity within the wealth distribution, particularly among unconstrained, where given the same level of liquid resources, households facing a lower level of income risk (*safe*) display larger MPCs, 0.50 on average, compared to households facing a higher level of income risk (*risky*), who exhibit an average MPC of 0.29. This result is robust to different measures of income risk and wealth⁵.
- (b) Theoretically, I show that the empirical findings can be easily rationalized by extending the (Kimball 1990b) two-period model with two consumers having the same characteristics found in the data for safe and risky households. Crucially, I show that the results hold given the same preferences and amount of optimal savings.
- (c) Numerically, the empirical results can be replicated almost perfectly by a quantitative heterogeneous agent standard incomplete market model (SIM)⁶ augmented with ex-ante heterogeneity in the income processes and calibrated to match the observed negative correlation between income level and income risk. Again, the results hold given the same discount factor and matching for both groups the same amount of wealth as in the data⁷.
- (d) The fiscal policy implications of the new source of MPC heterogeneity uncovered in this paper are far from negligible. I show that a stimulus check targeted to specific labor categories yields maximum impact on aggregate output only when targeted to civil servants, i.e., the safest category among employees. This fact implies that governments could face a trade-off between reducing the cost of labor income risk and boosting aggregate demand.

⁵I show that the results hold even by proxying income risk with labor categories, where Civil Servants - clearly a category facing lower uncertainty - exhibit higher MPC compared to private sector employees or self-employed households. Further, the results are verified regardless of the underlying measure of wealth: they hold comparing MPC within liquid assets only, within cash-on-hand - commonly defined as liquid assets plus monthly income- or within total wealth.

⁶Because these exercises are particularly relevant for current Heterogeneous Agent New Keynesian models (HANK), in which the MPC is a central concept (Kaplan and Violante 2021), I prove the results using the SIM model, which is the backbone of any HANK model.

⁷It should be noted that this is not a trivial result. Typically models with ex-ante heterogeneity in other dimensions feature a clear polarization of the wealth distribution, defined by (Kaplan and Violante 2021), as the "*missing middle*" problem. This can be the case, for example, by assuming ex-ante heterogeneity in discount factors as in (Aguilar, Bils, and Boar 2020)

To provide a clear understanding of the mechanisms at play, the simple two-period framework illustrated in Section 3, formalizes the empirical approach by studying the consumption and saving dynamics of two consumers: one with high income and low risk (safe) and another with low income and high risk (risky).

The intuition is the following: a lower income level implies that the risky consumer has a lower consumption level than the safe one, raising her marginal utility. Because she also faces more uncertainty, her expected marginal utility will be also higher. Therefore, both sides of the risky consumer's Euler equation are shifted up compared to those of the safe one. This fact allows both consumers to pick up the same amount of optimal savings. This is how the negative correlation between income levels and income risk creates wealth homogeneity⁸. Finally, note that heterogeneity in marginal utility also implies heterogeneity in the coefficient of absolute prudence, $-u'''(c)/u''(c)$, which in turn, implies a higher precautionary premium to compensate for the effect of income risk. At this stage, I show that given the same (or almost the same) wealth level, households with a higher precautionary premium exhibit a lower MPC.

Crucially, in all exercises, I show that when the negative correlation between income level and income risk is neglected, it is impossible to compare MPC within wealth and to isolate the effects of income risk on the MPC. This is because, given the same income level, when preferences are equal, households facing higher risk are much more wealthy than households facing lower risk. In contrast, given the same risk, households with a lower income level accumulate fewer assets than households with a higher income level. In both cases, it is impossible to mute the effect given by the concavity of the consumption function, which implies a lower MPC at higher levels of wealth (Carroll and Kimball 1996).

Therefore, one of the main contributions of this paper is to show that the effects of income risk on the MPC can be identified only through the interaction between income level and income risk and that ignoring this interaction leads to MPCs being driven primarily by wealth effects, masking the effects of income risk on the MPC⁹.

This result might explain why (Zeldes 1989,a) concludes that heterogeneity in income risk has unclear effects on MPC. In his paper, after documenting that the MPC under uncertainty is seven times the MPC under certainty, he also compares the MPCs among different income processes. However, he only assumes different variances of the income

⁸Clearly, in the data, or in models with full heterogeneity (as shown in Section 4), it is enough that the two masses of households have *almost* the same wealth or liquid assets-to-income ratio, i.e. absence of wealth polarization.

⁹This theory, might only stand if one compares households that potentially can have same preferences. That is, the effect of income risk on the MPC cannot easily be evaluated by comparing an entrepreneur with an employee.

processes.¹⁰

Finally, I show that in neither case the MPC depends on (permanent, or current) income levels, whose direct effect on the MPC is always ambiguous. This fact is hard to reconcile with empirical evidence provided by (Kueng 2018), who documents that the MPCs out of predetermined payments from the Alaska Permanent Fund are monotonically increasing in income. He proposes a behavioral explanation for his result. Similarly, in contemporary work, (Commault 2022a) documents empirically and theoretically that households with higher persistent income have higher MPCs out of transitory shocks. She motivates this result by showing that these households consume more, have a stronger precautionary motive, and their consumption is more responsive to transitory shocks. However, her empirical analysis does not identify this result within the wealth distribution. Instead, she shows that an increase in persistent earnings by one standard deviation raises the consumption response to transitory shocks by 6%-8%.

Therefore, my analysis suggests that there could be a possible omitted variable bias driving the empirical results of these papers, the correlation between income level and income risk. A further element that raises this suspect is given by the fact that although (Commault 2022a) 's theory assumes that income levels and income risk are not correlated, one of the key elements allowing her life-cycle model to reproduce the empirical results and provide a within wealth comparison is to account for the negative correlation between permanent income and unemployment risk¹¹.

Related Literature - This paper is connected with two strands of the literature. The primary connection is with the literature that studies the impact of non-capital income uncertainty on the MPC out of transitory shocks. The theoretical side of this literature, (Zeldes 1989,a; Kimball 1990a,b; Carroll and Kimball 1996), compares the MPC under uncertainty with the MPC under certainty equivalence. I use the theory developed by these papers and extend it to evaluate the effect of *heterogeneity* in labor income risk exposure on the MPC as described above¹².

¹⁰Should be noted that given the computational power available in the late '80s, (Zeldes 1989,a) could not compute the stationary distribution of wealth and elaborate the dynamics documented in this paper.

¹¹However, the models we consider are different, as she works with a rich life-cycle model. Moreover, her theory relies on the standardization of the entire consumption problem by permanent income as in (Carroll 2009). Something not feasible in my case, since I let the variances of shocks be a function of permanent income. Therefore, our respective numerical analyses might provide complementary evidence.

¹²Another related paper is (Luengo-Prado and Sørensen 2004). They simulate a buffer-stock model of consumption, explicitly aggregate over consumers, and estimate aggregate marginal propensities to consume out of current and lagged income using simulated data generated by the model. They also document that high uncertainty reduces the MPC. However, besides focusing on aggregate dynamics only,

The empirical side of this literature is quite scant, with few notable cases. (Friedman 1957), found that US farmers and self-employed households had lower MPCs. More recently, (Sala and Trivián 2021), using the Spanish Survey of Households Finance, found a lower MPC for households whose reference person is self-employed. As documented in Section 2, my results are in line with these papers. Finally, using Australian data, (Berger-Thomson, Chung, and Mckibbin 2009) reports that people with higher subjective labor income uncertainty tend to spend less out of additional income.

The second connection is with the broad empirical literature documenting heterogeneity in MPC out of transitory income changes. These papers, usually focus on four different dimensions: *i*) The role of household resources, (Jappelli and Pistaferri 2010b, 2014, 2020; Fisher et al. 2020; Bunn et al. 2018; Christelis et al. 2019) among others; *ii*) the type of transitory shock, highlighting the role of its sign and size, (Christelis et al. 2019; Fuster, Kaplan, and Zafar 2020; Bunn et al. 2018; Andreolli and Surico 2021); *iii*) the role of heterogeneity in households portfolio composition, distinguishing between liquid and illiquid assets holding, (Kaplan, Violante, and Weidner 2014); and *iv*) the role of debt and mortgage repayment, (Mian, Rao, and Sufi 2013; Misra and Surico 2014; Romer 2021; Kosar et al. 2023). These papers are also different according to the empirical approach used. I follow those papers that use semi-structural methods to estimate the MPC (Blundell, Pistaferri, and Preston 2008; Kaplan, Violante, and Weidner 2014; Commault 2022b; Ganong et al. 2020). Regardless of the specific approach and findings ¹³, these studies convey the fact that there is a significant degree of MPC heterogeneity along the wealth distribution. In this paper, I document that the MPC is heterogeneous also within the wealth distribution.

Outline - Section 1 presents the data. Section 2 documents individual labor income risk and presents the BPP estimator, showing the empirical results and robustness checks. Section 3 rationalizes the empirical evidence in a simple two-period model. Section 4 presents the full quantitative model. Section 5 studies the fiscal implications. Section 6 concludes.

they do not provide a clear identification strategy for quantifying the effect of labor income risk on the MPC, and do not document any clear relationship between income risk, income level, and MPC, as I do in this paper.

¹³See (Jappelli and Pistaferri 2020) and (Kaplan and Violante 2021) for a survey on methods and results.

1. Data

The empirical analysis of this paper is conducted mainly using panel data from the Bank of Italy SHIW, a representative sample of the Italian resident population, collecting detailed information on demographics, household consumption, income, and real and financial wealth¹⁴. Flow income and consumption data refer to the previous year; wealth and debt variables are end-of-period values. Questions concerning the whole household are answered by the head of the family or by the person most knowledgeable about the family finances. Starting from 1998, the SHIW is biannual.¹⁵ The last available survey year refers to 2020. I use the Historical Sample 1998-2016, excluding 2020 because of the four years jump. I will call this data *H-Sample* for future reference. The Cross-sectional sample size is, on average, around 7,700 households. The share of the panel component is, on average, 50%.

Measure of Income - The survey allows for alternative measures of income (from labor, capital, transfers), and households report their net income. Consistent with the interest of detecting the direct impact of non-capital income uncertainty on the MPC, I follow (Kaplan, Violante, and Weidner 2014) and consider labor income plus government transfers. This measure represents a regular inflow of liquid wealth. I exclude interests, dividends, and other capital income because they are realized more infrequently.

Measure of Liquid Wealth - This is given the sum of checking, saving, money market, and call accounts plus directly held mutual funds, stocks, corporate bonds, and government bonds net of consumer debt.

Measure of Illiquid Wealth - The definition closely follows (Kaplan, Violante, and Weidner 2014). Therefore, Illiquid wealth is the value of the household's property net of mortgages and other loans to purchase the home, retirement plans, cash value of life insurance policies, certificate of deposits, and saving bonds. All monetary variables are expressed in the 2016 euro using the CPI.

Measure of Consumption -, I adopt the standard measure of non-durable consumption, following (Blundell, Pistaferri, and Preston 2008; Kaplan, Violante, and Weidner 2014; Jappelli and Pistaferri 2020), This measure includes food at home and food away from home, utilities, gasoline, car maintenance, public transportation, child care, health expenditures, and education.

¹⁴The Historical dataset (as well as annual waves) is publicly available with documentation in English at the Bank of Italy's website, see: <https://www.bancaditalia.it/statistiche/tematiche/indagini-famiglie-impres/bilanci-famiglie/distribuzione-microdati/index>

¹⁵The SHIW has been conducted annually from 1965 to 1987 (except for 1985). Since 1987 the survey has been conducted every other year except for a three-year interval between 1995 and 1998. More accurate information on asset holdings is collected from 1987 after including wealthier households.

Sample Selection - I start from the H-Sample 1998-2016 and drop households with intermittent "headship" or those with missing information on education or region of residence. I keep only households where the head is 25-55 years old. This is done to minimize the number of people with lower attachment to the labor force. I keep only households that never moved from south to north, or vice-versa. This choice is important since there is considerable economic divergence among these regions in Italy.

TABLE 1. - SUMMARY STATISTICS, SHIW: H-SAMPLE 1998-2016

Demographics			Financials ($\times 1.000$)		
	Mean	SD		Mean	SD
Male	0.60	0.49	Disposable labor income	26.01	14.55
Age	43.28	7.08	Net liquid wealth	15.24	48.97
Married	0.74	0.44	Net illiquid wealth	107.13	143.34
Years of education	7.17	4.96	Net Worth	121.41	157.87
Resident in the South	0.37	0.48	Non-durable consumption	22.05	9.73
Size of the family	3.18	1.30			
Number of labor income earners	1.47	0.71			
City size >500,000	0.08	0.28			
Number of observations	7,341.00	0.00			

Statistics after sample selection.

Concerning income, I drop households whose income grows more than 500%, falls by more than 80%, or is below 100 Euro, and trim the 99th income percentile. I drop households with negative or zero consumption levels. I keep only households that appear in the sample at least three consecutive times because identifying the coefficients of interest requires a minimum of three periods¹⁶.

The strong and conservative sample selection ensures the well-known measurement error in survey data, particularly for consumption expenditures, to be minimal. However, it comes at the cost of reducing the available sample size and potentially underestimating income risk. The final sample has 7.341 observations over the pooled years 1998-2016. Table 1 reports basic statistics for the H-sample 1998-2016. Detailed statistics for safe and risky households are presented later¹⁷.

¹⁶All these requirements would not be necessary to estimate the CV. For instance, (Arellano et al. 2022), also consider households with zero income.

¹⁷Potentially, the SHIW contains another well-known unique resource, a panel of self-reported MPCs, dataset already used in (Jappelli and Pistaferri 2010b, 2014, 2020). I tried to use it, particularly for testing the relationship between the change in income risk and the change in self-reported MPCs. The results show some negative correlation. However, the small sample size prevents making further conclusions. These results are available upon request.

Finally, for external validity purposes and comparability with preexisting studies, I also consider the PSID sample 1999-2010, already used in (Kaplan, Violante, and Weidner 2014).¹⁸ I will call this data *US-Sample* for future reference.

2. Empirical analysis

This section proceeds in three steps. First, I compute the Income risk and identify risky and safe households. Second, I compute both groups' MPC out of transitory shocks and test eventual differences. Finally, I consider some robustness checks. Additional estimations and robustness checks are relegated to the Online Appendix.

2.1. Measuring Income Risk

To measure income risk, I follow the methodology proposed by (Arellano et al. 2022), defining risk as the part of income changes that the households cannot predict. This concept differs from income volatility or instability, as in (Haider 2001; Gottschalk and Moffitt 2009; Hoffmann, Malacrino, and Pistaferri 2021). It implies defining income risk using expected income values based on critical determinants of the household's information set. Accordingly, I construct a coefficient of variation (CV) measure of income risk as a ratio between two measures, the mean absolute deviation, which is a measure of the dispersion of the predictive distribution of income, and the mean, which is a measure of location:

$$(1) \quad CV(X_{i,t}) = \frac{\mathbb{E}(|Y_{i,t} - \mathbb{E}(Y_{i,t}|X_{i,t})||X_{i,t})}{\mathbb{E}(Y_{i,t}|X_{i,t})}$$

The CV is an intuitive measure of the total uncertainty households face. For instance, an individual with an expected income of 20,000 euros and a CV of 0.1 faces a deviation of her next year's income from its mean of ± 2000 euros. Using the mean absolute deviation instead of the standard deviation in the numerator aims to minimize the sensitivity to extreme observations. The set of predictors $X_{i,t}$ is supposed to exhaust the information set of a household. It includes micro and macro predictors at the national and regional levels, reported in Table 2. All predictors have further interacted with a quadratic in age, year, and macro-area dummies. The distribution of income $Y_{i,t}$ is in level, given predictors $X_{i,t}$. Lagged income values are entered in logarithmic form. Since income is positive, the

¹⁸This dataset is publicly available on Greg Kaplan's website <https://gregkaplan.me/academic-publications>

parametric estimator is recovered from the two following exponential specifications:

$$(2) \quad \mathbb{E}(Y_{i,t}|X_{i,t}) = \exp(X'_{i,t}\beta)$$

$$(3) \quad \mathbb{E}(|Y_{i,t} - \mathbb{E}(Y_{i,t}|X_{i,t})||X_{i,t}) = \exp(X'_{i,t}\gamma)$$

TABLE 2. - CV PREDICTORS, SHIW: H-SAMPLE 1998-2016

Micro		Macro
Income	Demographics	National, Regional
Log lagged income	Age	Unemployment rate (U)
Log lagged income squared	Age squared	GDP
Log lagged income cubed	Edu: lower secondary	$U \times$ South Regions
Indicator lagged income = 0	Edu: upper secondary	GDP \times South Regions
Days worked (t_{-1})	Edu: college and higher	
Indicator out-of-work income = 0	Regional dummies	
Indicator full time ($t, t_{-1}, t_{-2}, t_{-3}$)	Number of Income gainers	
Extra-hours worked (t, t_{-1})		
Permanent worker		
Firm size		
Industrial Sector		

Where γ and β are estimated using two separate Poisson regressions¹⁹.

In practice, the estimation requires two-stage, first estimates $Y_{i,t}$ on $X_{i,t}$, which gives the CV denominator, then taking the residual of the latter and further regressing it on $X_{i,t}$, one obtains the numerator.

The key advantage of this methodology is that including a rich set of variables in the prediction step allows capturing the dispersion that is less expected, will occur tomorrow, and (hopefully) last permanently. To see this point is useful to consider the non-parametric measures that use differences in income to document income risk, (Güvenen, Ozkan, and Song 2014). The main drawback of this approach is that while taking differences in income, one does not know if the income change was predictable. The role of covariates in predicting the CV is crucial to remove the predictable part and isolate the change that was not predicted.

¹⁹Poisson regression is numerically more stable than exponential regression, especially when there are small sample size concerns.

Further, as with other measures, also the CV is time variant. Therefore, one can eliminate possible sources of unobserved heterogeneity, like risk aversion or preferences for particular types of professions.

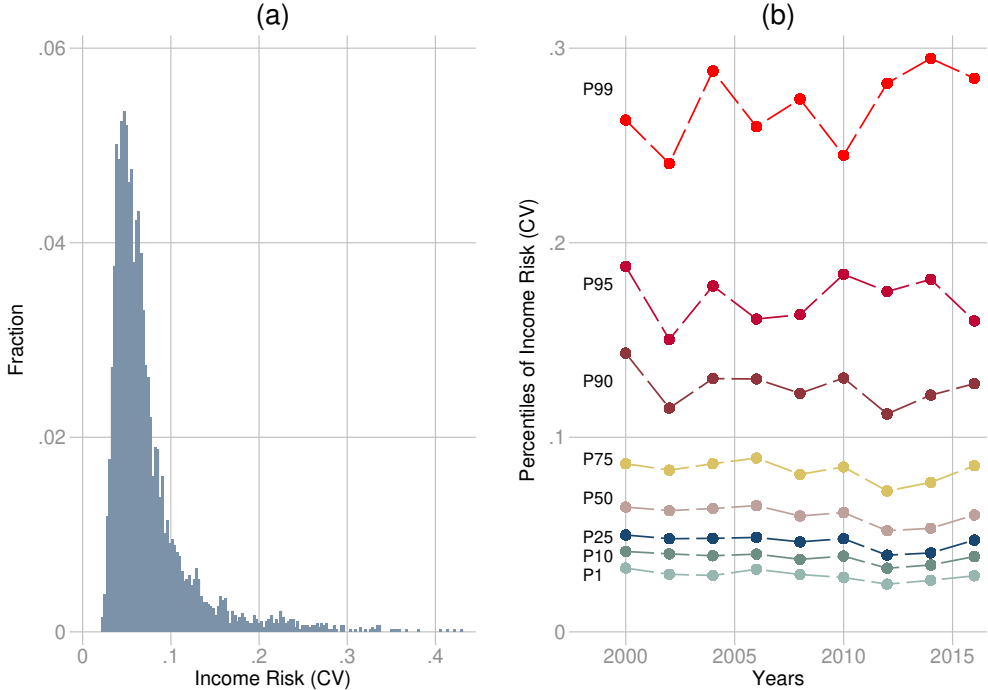


FIGURE 1. Labor Income Risk in Italy
 Author’s calculations. SHIW: H-SAMPLE 1998-2016

Unfortunately, also the CV has its limitations: this is a measure of total income risk, and no clear procedures exist to identify the *expected* transitory and permanent components separately. In principle, it can be shown that using more lags to predict future income implies incorporating more permanent income than a transitory one. However, this approach would considerably reduce the number of observations without large data at hand. Section 4 deals with this issue.

It is also important to stress that conditional higher-order income moments such as skewness and kurtosis also matter when assessing income risk, (Güvenen, Ozkan, and Song 2014). As a result, the CV will possibly underestimate the uncertainty individuals face. I assess this point in the online Appendices B5 and B6, where I compute the kurtosis and the Kelley skewness using quantile regressions.

Income Risk in Italy and the identification of Safe and Risky households

Figure 1 reports the distribution of income risk and various percentiles of the distribution over time. Two striking facts emerge: First, income risk, as measured from the SHIW, is not high; the average CV is around 0.07. Second, half of the Italian economy faces low uncertainty, which seems stable over time. Accordingly, I identify those households with a CV below or equal to the median CV as *safe*. Otherwise, households are classified as *risky*.

In documenting the CV, it is useful to use as a benchmark the results obtained by other studies that compute the CV using administrative data, (Arellano et al. 2022) and (Lagrosa 2022).

A reassuring fact comes when comparing my estimates with those provided by (Lagrosa 2022). He uses an administrative longitudinal random sample compiled by the Italian National Social Security Agency (INPS). Qualitatively, the trends are similar, and he also documents that half of Italian households face low-income risk. This fact confirms the representativeness of the SHIW²⁰. Quantitatively, I document slightly lower values; his average CV is 0.12, while, as aforementioned, I report an average of 0.07. Although the difference is small, this fact is not surprising. I deal with a smaller sample size, a strong sample selection, and the biannual nature of the survey. Further, even if the methodology used is the same, the results change also according to the type and the quality of predictors used. I directly control for variables like the number of income gainers in the family and extra hours worked in past years. These variables could account for changes in labor supply to cope against risk, lowering the quantitative risk that the households cannot predict. Finally, while he does not analyze the correlation between the CV and income levels, we provide similar numbers for the average CV of permanent workers, around 0.04.²¹

(Arellano et al. 2022) use high-quality administrative data from Spain. Still, the comparison makes sense from a qualitative point of view as Spain and Italy share common micro and macroeconomic features. These authors document four important results from their analysis: first, income risk is inversely related to income and age. Second, income risk inequality increases markedly in periods of recessions. Third, they found little variability between years, suggesting persistence in income risk inequality. Finally, the quantitative role of aggregate conditions in explaining the evolution of individual income risk is small.

²⁰As shown in the Online Appendix B1, I also show that the evolution of upper and lower inequality, based on the CV, is similar to what is shown by (Hoffmann, Malacrino, and Pistaferri 2021). These authors use Italian administrative data from 1980 to 2020 and quantify income risk using a volatility measure, which should be somehow correlated with the CV. Dynamics from 1998 to 2016 provides another reassuring comparison.

²¹I cannot provide consistent estimates for temporary workers as he did: After sample selection, the number of temporary workers is limited, representing a direct source of underestimation in my case.

At the same time, a strong predictor is given by the number of days/months worked in the past year. All results that, at least from a qualitative point of view, I confirm in this paper. Quantitatively, since my results are closer to (Lagrosa 2022), while (Arellano et al. 2022) reports a higher CV, this suggests that Italian households face lower income risk than the Spanish one. This fact is unsurprising given that more than 1/3 of the Spanish workforce has temporary contracts.

TABLE 3. - SUMMARY STATISTICS BY RISK GROUP, SHIW: H-SAMPLE 1998-2016.

Statistics	Risky Households		Safe Households	
	Mean	SD	Mean	SD
<i>Demographics</i>				
Male	0.58	0.49	0.63	0.48
Age	42.37	8.52	45.71	3.91
Married	0.70	0.46	0.80	0.40
Years of education	7.24	4.96	7.00	4.71
Resident in the South	0.43	0.49	0.28	0.45
Size of the family	2.97	1.33	3.38	1.23
Number of labor income earners	1.35	0.65	1.70	0.68
City size >500,000	0.09	0.28	0.07	0.25
<i>Financials</i> ($\times 1.000$)				
Disposable labor income	22.12	12.22	30.68	14.26
Net liquid wealth	12.69	55.87	17.50	52.36
Net illiquid wealth	87.62	123.38	117.64	135.55
Net Worth	99.53	139.18	134.06	154.04
Non-durable consumption	19.88	8.34	24.83	10.08
<i>Income Risk</i> (2000-2016)				
Coefficient of Variation (CV)	0.102	0.05	0.046	0.01
Number of observations	2,314.00	-	2,318.00	-

Statistics after sample selection.
Risky: $CV > p50(CV)$, Safe: $CV \leq p50(CV)$

Correlates of Income Risk - Before computing the MPC, it is instructive to consider some correlates of income risk, the distribution of liquid assets on a logarithmic scale, and some descriptive statistics for both safe and risky households.

As shown from panel A in figure 2, there is a strong negative correlation between income level and income risk²². Panel B shows more dispersion in the correlation between income risk and net-liquid assets. This fact is consistent with different strands of the

²²As shown in the Online Appendix A, Figure A3 - panels C and D, this is true for lagged and permanent income.

literature, pointing out that higher-income risk leads to higher liquidity holdings,²³ or the portfolio of a rational household with higher labor income risk can be well-diversified (Guiso, Jappelli, and Terlizzese 1996). However, the fact that low-income households face higher income risk prevents a positive correlation. As a result, the relationship between income risk and liquid assets is almost flat.



FIGURE 2. Correlates of Income Risk and Liquid Assets by groups.

Author's calculations. SHIW: H-SAMPLE 1998-2016

The figures include demographic variables and year-fixed effects.

Risky: $CV > p50(CV)$, Safe: $CV \leq p50(CV)$

This fact is reflected in panel D, showing the distribution of log net-liquid assets for both safe and risky households, indicating small average differences between the groups, as documented in Table 3. The dynamics behind this fact are further rationalized in Sections 3 and 4. What matters here is the absence of polarization in net-liquid asset holdings, as this allows me to compare the estimated MPCs within the distribution itself. Finally, also from panel C emerges an almost flat and more dispersed relationship between income risk

²³See (Pratt and Zeckhauser 1987; Guiso, Jappelli, and Terlizzese 1992; Bertaut and Haliassos 1997; Koo 1998; Heaton and Lucas 2000; Bayer et al. 2019)

and net-illiquid assets. This fact is not surprising since the 77% of Italian households own at least one house.²⁴

I am now in a position where I can compute the MPC for both groups. Further details on the CV, its determinants, and covariates are documented in the Online Appendix A.

2.2. MPC Estimation

This section summarizes the construction of the (Blundell, Pistaferri, and Preston 2008) estimator for the MPC. The *BPP methodology* is a two-stage procedure: first, I regress log income and log consumption expenditures on year and cohort dummies, education, family structure, employment, geographic variables, and interactions of year dummies with education, employment, and region. Second, I construct the first-differences residuals of log consumption $\Delta c_{i,t}$ and log income $\Delta y_{i,t}$. Following (Kaplan and Violante 2010), the income process is represented as an error component model comprising orthogonal permanent and i.i.d. components. This formulation has been proved to be robust by (Commault 2022b) for biannual data²⁵. Using the Income growth formulation, one gets the well-known process presented in (MaCurdy 1982; Abowd and Card 1989)

$$(4) \quad \Delta y_{i,t} = \eta_{i,t} + \Delta \varepsilon_{i,t}$$

The modified BPP estimator of the transmission coefficient of transitory income shocks to consumption (the MPC), proposed in this paper is simply the standard BPP estimator conditional on the CV, given by:

$$(5) \quad \widehat{MPC}_{j,t} = \underbrace{\frac{\text{cov}(\Delta c_{i,t}, \Delta y_{i,t+1})}{\text{cov}(\Delta y_{i,t}, \Delta y_{i,t+1})}}_{\text{Realized income flows}} \bigg| \underbrace{\frac{\mathbb{E}(|Y_{i,t} - \mathbb{E}(Y_{i,t}|X_{i,t})||X_{i,t})}{\mathbb{E}(Y_{i,t}|X_{i,t})}}_{\text{CV}_{j,t}, \text{ Less Expected Income flows}}$$

Where:

$$j = \begin{cases} \text{Safe} & \text{if } CV \leq P50_{cv} \\ \text{Rrisky} & \text{if } CV > P50_{cv} \end{cases}$$

²⁴The 2019 census published by the Italian Ministry of Economy and the Tax Authority (Agenzia Delle Entrate) reveals that 77% of Italian people live in a property purchased or inherited. Only 23% of Italians live in rent: mostly young families or families residing in the north where housing prices are higher.

²⁵(Commault 2022b) shows that if transitory shocks are persistent, the estimates are biased.

The true marginal propensity to consume out of a transitory shock is defined as:

$$(6) \quad MPC_{j,t} = \frac{cov(\Delta c_{i,t}, \varepsilon_{i,t})}{var(\varepsilon_{i,t})} | CV_{j,t}$$

To implement the BPP, two identifying assumptions are needed: *i*) households have no foresight about permanent future shocks; *ii*) households have no advanced information about transitory future shocks; That is:

$$(7) \quad cov(\Delta c_{i,t}, \eta_{i,t+1}) = cov(\Delta c_{i,t}, \varepsilon_{i,t+1}) = 0$$

The advanced information/foresight issue in BPP is that households might adjust consumption today in anticipation of a permanent/transitory shock tomorrow.

To this end, it is crucial to stress that using the CV does not violate the identification assumptions. The reason is twofold. First, the BPP is computed on realized income values. Second, the CV, by predicting income (hopefully predicting the permanent shock that occurs tomorrow), is aimed at removing the predictable part and providing the change that was not predicted by the household. Recall that given the rich set of variables included in the prediction step of the CV, it is capturing dispersion that is less expected.

The final estimator is implemented by an IV regression of $\Delta c_{i,t}$ on $\Delta y_{i,t}$, instrumented by $\Delta y_{i,t+1}$ conditional on $CV_{j,t}$. Note that $\Delta y_{i,t+1}$ is correlated with the transitory shock at time t but not with the permanent one.

Moreover, (Kaplan and Violante 2010) provides an important result, showing that tight borrowing constraints do not bias the transmission coefficient estimate for transitory shocks.

2.3. Results

This subsection presents the MPCs for both groups of households, safe and risky. Recall that the survey is biannual. Therefore the MPC represents how much the households are going to spend within two years after the shocks.

In Table 4, I compute the MPC for safe and risky households conditional on the level of net-liquid assets. I further split the sample between low and high liquid assets (below or above the median of liquid assets)²⁶. This baseline specification suggests that either below or above the median of the liquid asset distribution Safe households have a statistically

²⁶The small sample size prevents further fragmentation. However, the online appendix B.7 shows that the results are confirmed even by splitting the sample into quartiles.

significant higher MPC.

Interestingly, when looking above the median of liquid assets holdings, results indicate that almost 25% of the unconstrained households in Italy not only face low levels of income risk but also exhibit very high MPCs out of transitory shocks²⁷.

TABLE 4. BIANNUAL MPC OUT OF TRANSITORY INCOME SHOCKS: ABOVE AND BELOW MEDIAN INCOME RISK, AND LIQUID ASSETS HOLDINGS

	Low Liquid Assets		High Liquid Assets	
	Risky	Safe	Risky	Safe
MPC	0.240*** (0.062)	0.469*** (0.050)	0.338*** (0.053)	0.538*** (0.051)
<i>p-values of equality test</i>		0.01		0.01
Observations	865	757	665	933

H-Sample 1998-2016

Risky = Above median Income risk, Safe= Below median Income Risk.

Low Liquid Assets = Below median Liquid Assets

Bootstrapped standard errors based on 250 replications in parenthesis.

To test whether these results survive when accounting for the total portfolio composition, I follow (Kaplan, Violante, and Weidner 2014) and split the sample between Wealthy-Hand-to-Mouth - those households with a liquid wealth of less than half their monthly income and sizable positive amounts of illiquid assets -, and Non-Hand-to-Mouth, those households who are essentially unconstrained. The share of Poor-Hand-to-Mouth, those households whose liquid and illiquid assets level, is meager in Italy, around 5%. Therefore, I estimate the effect only for Wealthy-Hand-to-Mouth (19% of the final sample), and Non-hand-to-mouth (76% of the sample). Similar shares for Italy have been documented by (Kaplan, Violante, and Weidner 2014). Using the European Central Bank HFCS²⁸, these authors find an average share of Poor-Hand-to-Mouth of around 8.3% and an average share of Wealthy Hand-to-Mouth equal to 15.5%.²⁹

²⁷Recall that the BPP methodology provides average point estimates of the MPC. Therefore, these results do not necessarily imply a positive relationship between MPC and liquid assets. To see this point, consider that, for both groups of households, the MPC differential between low and high liquid assets (0.240 vs 0.338 for risky and 0.469 vs 0.538 for risky) are not statistically different from zero. Indeed, using fractional polynomial regressions, the online appendix B.8 (Figure A4) shows a negative relationship between the estimated MPC and cash on hand for both groups of households. This negative relationship is consistent with (Jappelli and Pistaferri 2020). In this case, I show that the MPCs for risky households decrease faster over the liquid assets/cash-on-hand distribution.

²⁸HFCS = Household Finance and Consumption Survey.

²⁹Since I closely follow (Kaplan, Violante, and Weidner 2014) for the computation of these classifications,

In contrast with the results provided in Table 4, for households with few liquid asset holdings, the MPC of the Wealthy-Hand-to-Mouth households is unaffected by the level of income risk. While, in line with previous results, the SHIW uncovers the presence of the "Safe"-Non-Hand-to-Mouth households. This wealthy fraction represents almost 38% of the Italian population and strongly responds to transitory income shocks.

TABLE 5. BIENNIAL MPC OUT OF TRANSITORY INCOME SHOCKS: ABOVE AND BELOW MEDIAN INCOME RISK, HAND TO MOUTH STATUS

	W-Hand-to-Mouth		N-Hand-to-Mouth	
	Risky	Safe	Risky	Safe
MPC	0.127	0.321***	0.330***	0.539***
	(0.117)	(0.097)	(0.043)	(0.040)
<i>p-values of equality test</i>		0.19		0.00
Observations	312	270	1134	1373

H-Sample 1998-2016

Risky = Above median Income risk, Safe= Below median Income Risk.

Bootstrapped standard errors based on 250 replications in parenthesis.

Although the inconsistency occurring for low liquid assets households can be driven by measurement errors, the lack of significance in MPC differentials between the two groups of Wealthy-Hand-to-Mouth is not surprising and, as shown in the Online Appendix B.7, derives from a loss of differences in MPC among households close to the borrowing constraint. In other words, having more or less income risk might not be of first-order importance whenever a household is constrained, simply because it cannot adjust consumption and saving decisions according to the risk it faces. This result provides empirical support to the theoretical findings by (Carroll, Holm, and Kimball 2021). Indeed these authors show that adding income risk to liquidity constraints does not produce sizable effects on the MPCs.

Intuitively, from the perspective of an unconstrained household whose income stream is predictable, there are few incentives to adjust saving after a small windfall. In contrast, when the uncertainty band is larger (risky), and households can adjust their consumption path, the income risk effect dominates. Therefore when households are far away from the borrowing constraint, the higher the risk, the smaller the part of the transitory shock that is transmitted to consumption. Sections 3 and 4 try to formalize this argument while the full procedure and detailed sample statistics are presented in the Online Appendix D.

rationalizing this evidence.

Robustness - A natural check is whether using a different proxy for labor income risk would yield the same results. Although the SHIW contains measures of subjective expectations, I do not make use of them for two reasons: first, the BPP requires that households must be observed at least for three consecutive periods to remove unobserved heterogeneity. During the sample 1998-2016, there are no three consecutive questions on subjective risk expectations preventing this application in this context. Second, there is a much more intuitive way of inferring labor income risk: using labor categories. Table 6 compares - within wealth, as before - the MPC for civil servants, the safest labor category³⁰, against the MPC deriving from other employees working in the private sector and self-employed households³¹. Interestingly, the income of civil servants is also higher than the income of private sector and self-employed households.

TABLE 6. BIENNIAL MPC OUT OF TRANSITORY INCOME SHOCKS: LABOR CATEGORIES, BELOW AND ABOVE MEDIAN LIQUID ASSETS HOLDINGS.

	Low Liquid Assets			High Liquid Assets		
	Other	Civil	Self	Other	Civil	Self
MPC	0.340*** (0.040)	0.433*** (0.084)	0.274 (0.419)	0.317*** (0.038)	0.512*** (0.074)	0.225* (0.111)
<i>p-values of equality test</i>		0.31	0.68		0.02	0.02
Observations	1587	355	159	1445	453	264

Other = Households whose reference person is employed in the private sector or is self-employed. Bootstrapped standard errors based on 250 replications in parenthesis. p-values of equality test compare the MPC of other vs. Civil servants and Self-employed vs. Civil servants.

The results confirm the evidence that once wealth is fixed, households with higher income and low levels of income risk (Civil Servants) have higher MPC. The lower MPC for self-employed households, although probably driven by measurement errors given by the small sample size, is in line with (Friedman 1957), documenting that US farmers and self-employed households had lower MPCs. Further, (Sala and Trivín 2021), using the Spanish Survey of Households Finance, found a lower MPC for households whose reference person is self-employed.

³⁰As a reference point, civil servants have a CV value ranging between 0.02 and 0.03, which is between the 1st and 10th income risk percentiles; private sector employees display a CV that is between the 10th and the 75th percentile, Self-Employed households have a CV comparable to those of employed households with a CV from the 25th to the 99th percentile.

³¹These should not be interpreted as entrepreneurs.

Finally, given that households with low risk have higher income, in Table 7, I explicitly test if these results are driven by income. Because the CV negatively correlates with current, lagged, and permanent income³², I repeat the same exercises as before, replacing the CV with one of the three measures of income. Therefore, I create two groups: low and high income (above median income) and test the differences below and above median liquid assets.

TABLE 7. BIENNIAL MPC OUT OF TRANSITORY INCOME SHOCKS: ABOVE AND BELOW MEDIAN INCOME LEVEL, AND LIQUID ASSETS HOLDINGS

	Low Liquid Assets		High Liquid Assets	
	Low Income	High Income	Low Income	High Income
<i>Current Income</i>				
MPC	0.312***	0.424***	0.305***	0.372***
	(0.050)	(0.101)	(0.058)	(0.061)
<i>p-values of equality test</i>		0.40		0.90
<i>Lagged Income</i>				
MPC	0.355***	0.357***	0.400***	0.334***
	(0.047)	(0.060)	(0.050)	(0.048)
<i>p-values of equality test</i>		0.92		0.30
<i>Permanent Income</i>				
MPC	0.359***	0.344***	0.388***	0.344***
	(0.047)	(0.061)	(0.055)	(0.047)
<i>p-values of equality test</i>		0.84		0.43

SHIW 1998-2016

Bootstrapped standard errors based on 250 replications in parenthesis.

In neither case the MPC is statistically different from zero. This means that once wealth is fixed, income does not play a direct role in determining the MPC. This fact suggests that households with low income risk have higher MPC. Therefore, papers documenting

³²Permanent income is constructed as the average income over the entire period until the households exit.

that high-income households have higher MPC, (Kueng 2018; Boutros 2021; Commault 2022a), which usually control for wealth rather than testing differences within-wealth, can be subject to a bias given by the missing interaction between income levels and income risk.³³

2.4. Additional Results

Additional robustness The Online Appendix B contains various robustness exercises. First, I show that the results are unaffected when estimations only include male households; if the definition of liquid assets also includes monthly income (i.e. cash-on-hand); if non-durable consumption is replaced with total consumption. Then, I show that the results are consistent even using higher-order moments of income risk, like Kurtosis and Kelley Skewness. Finally, to explicitly account for other sources of unobserved heterogeneity, I show that even taking out first differences in income risk or liquid asset holdings, results are unaffected.

External Validity The Online Appendix E2, shows that applying the CV and same estimation procedure to the same PSID sample used in (Kaplan, Violante, and Weidner 2014), the Safe Non-Hand-to-Mouth exhibits a higher MPC than the Risky Non-Hand-to-Mouth. Results are particularly significant for male households.

3. Understanding the results extending the Kimball (1990b)'s two-period Model

Despite numerous endeavors to mitigate measurement errors, the empirical analysis continues to be susceptible to potential biases arising from small sample sizes. Consequently, this Section strives to methodically rationalize and formalize the empirical facts that have been presented thus far. The starting point is the two-period model presented by (Kimball 1990a). Consistently with the analysis made so far, I extend this model by considering two consumers: one, called safe (high labor income and low second-period uncertainty), and another, called risky (low labor income and high second-period uncertainty). Both consumers cannot borrow. The interest rate is exogenous and, without any loss of generality, is set to zero. The consumers' problem is:

³³However, this is not necessarily the case if there is committed consumption as shown by (Shore and Sinai 2010). If committed consumption is high, high-income households may have larger MPC because they allocate a higher fraction of their income in mortgages and/or installment expenditures, (Misra and Surico 2014; Kosar et al. 2023).

$$(8) \quad \max_{c_j} u(c_j) + \mathbf{E}v_j(w_{0,j} - c_j + y_j)$$

The subscript j stands for $s =$ safe consumer, $r =$ risky consumer. The first-period utility function is u , and c_j is the first-period consumption. \mathbf{E} an expectation conditional on first-period information, v is the second-period utility function. The initial cash on hand is given by w_0 and includes initial assets a_0 , plus the first-period labor income received before any decision occurs.

I assume that w_0 is the same for both consumers. The second-period labor income y_j , is defined as $y_j = \mu_{y,j} + \varepsilon_j$. Here, $\mu_{y,j}$ is the mean of second-period income (its expectation) and ε_j a mean zero income risk component. Consistently with the evidence from the data, I assume that the "safe" consumer earns more than the risky one, $\mu_{y,s} > \mu_{y,r}$, and also has a lower income risk, $\sigma_{\varepsilon,s}^2 < \sigma_{\varepsilon,r}^2$, such that, $y_s > y_r$. Note that the two respective risk levels are independent of each other.

Precisely as in (Kimball 1990b), I define $w_j = w_0 + \mu_{y,j}$ referring to the total human and non-human wealth and $s_j = w_j - c_j$ as the amount of saving out of total human and non-human wealth. In this context, the fact that $w_s > w_r$, from $\mu_{y,s} > \mu_{y,r}$, can be seen as a permanent heterogeneity like a skill premium acquired by the safe consumer.

Finally, the second period utility function $v(s_j)$ is of the HARA class, such that $v'(s_j) > 0$, $v''(s_j) < 0$, $v'''(s_j) > 0$ and absolute prudence, $-\frac{v'''(s_j)}{v''(s_j)} = \eta(s_j)$ is decreasing in s_j , through c_j . Both consumers have the same preferences.

Given this setting, the problem, as defined in equation 8, can be rewritten as:

$$(9) \quad \max_{c_j} u(c_j) + \mathbf{E}v_j(w_j - c_j + \varepsilon_j)$$

The first order condition yields:

$$(10) \quad u'(c_j) + \mathbf{E}v'_j(w_j - c_j + \varepsilon_j)$$

Therefore, uncertainty about second-period income affects consumption in the first period as long as it affects the second-period expected marginal utility.

In what follows, I will discipline the analysis to be as consistent as possible with what was discussed about the within-wealth MPC heterogeneity. In the original problem exposed by (Kimball 1990b), the analysis was to understand the difference between the MPC under

uncertainty against the MPC under the certainty equivalent case, and the comparison was at each level of consumption. In what follows, with a slight abuse of the information acquired in the data, I will compare the MPC at each wealth level.

3.1. Saving dynamics among the two consumers

As pointed out in the empirical section, the higher the precautionary motive, the higher the level of saving a consumer will find optimal to choose. Then, the concavity of the consumption function implies that the MPC decreases in wealth, (Carroll and Kimball 1996). This is a clear endogeneity problem that masks the pure effect of income risk on the MPC. Therefore, this subsection studies the dynamics of saving that allow both consumers to pick up the same optimal level of saving in equilibrium. This would (virtually) ensure that when comparing the MPCs between the two consumers, the wealth effects are somehow muted.

Proposition 1 - Compensating Precautionary Premium: Given two consumers $j = s, r$, here defined as $s =$ safe consumer, $r =$ risky consumer, having total wealth given by $w_j = w_0 + \mu_{y,j}$, non-capital income given by $y_j = \mu_{y,j} + \varepsilon_j$, where w_0 is the same for both and $\mu_{y,s} > \mu_{y,r}$, and $\sigma_{\varepsilon,s}^2 < \sigma_{\varepsilon,r}^2$, there exist some quantity ψ_j^* , increasing in $\sigma_{\varepsilon,j}^2$ such that can compensate for the effect of the risk on their respective second-period expected marginal utility. Therefore, first-period consumption would be unaltered by income risk. This quantity is such that:

$$(11) \quad v'_j(w_j - c_j) = \mathbf{E}v'(w_j - c_j + \varepsilon_j + \psi_j^*)$$

where:

$$(12) \quad \psi_j^*(s_j, \varepsilon_j) = \frac{-v_r'''(s_j) \sigma_j^2 \varepsilon}{v_r''(s_j) 2} + 0(\sigma_j^2 \varepsilon)$$

where $0(\sigma_j^2 \varepsilon)$ is a quantity that converges to zero faster than $0(\sigma_j^2 \varepsilon)$ as $\sigma_j^2 \varepsilon$ goes to zero. Proposition 1 establishes the presence of a *compensating precautionary premium*, that is, the level of wealth needed to compensate for the effect of the risk on consumption, or equivalently, the rightward shift in first-period consumption, needed to self-insure against

income risk. The higher the risk, the higher the wealth needed to cover it. ³⁴

Proposition 2 - Saving Homogeneity across groups: Given proposition 1, and provided that w_0 is the same for both and $\mu_{y,s} > \mu_{y,r}$, there exists a certain $cov(\mu_{y,j}, \sigma_{\varepsilon,j}) < 0$ able to generate a pair of consumption levels $\{c_s, c_r\}$ satisfying the optimal policy rule $c_j = w_j - s_j^*$ such that both consumers have the same amount of saving $s_j = w_j - c_j$ out of total human and non-human wealth. That is, $\underbrace{w_{0,s} + \mu_{y,s}}_{w_s} - c_s = \underbrace{w_{0,r} + \mu_{y,r}}_{w_r} - c_r$ if and only if $c_s > c_r$. Given the assumed initial conditions, the risky consumer must consume less to have the same amount of savings as the safe consumer.

Proposition 1 and 2 together illustrate the importance of the negative correlation between income level and income risk to generate the same amount of saving needed to compare the respective marginal propensities to consume. Note that $v_j(s)'' < 0$ implies that the left-hand side (LHS) of the Euler equation is increasing in s , while the right-hand side (RHS) will be decreasing. Therefore, there exists a level s_j^* uniquely determined for both consumer's problems.

The fact that the risky consumer has a lower initial μ_y implies that her LHS will be shifted up with respect to the safe consumer. That is, her marginal utility will be higher. This is illustrated in Figure 3. Furthermore, the fact that she faces more risk than the safe consumer means a higher compensating precautionary premium is needed. Therefore, her expected marginal utility will also be higher. Thus, there will be some pair c_j such that both consumers have the same level of savings, as depicted in Figure 3.

Given the simple problem studied here, the condition $s_s = s_r$ is more stringent than what can happen in the data or in a model with full heterogeneity. In these cases, this condition reduces to one in which there are sufficient masses of both types for each point of the wealth distribution, i.e., no wealth polarization. Still, this simple analysis allows us to rationalize the fact that in the data, the distribution of liquid assets among safe and risky households is not polarized, highlighting the role of the negative correlation between income levels and income risk.

Of course, the attentive reader may ask the following: what if the income risk of the

³⁴(Kimball 1990b), also makes clear the mutual inverse of the compensating precautionary premium. That is the equivalent precautionary premium ψ defined as the leftward shift of the consumption function that would result from eliminating the risk. This implies that equation 11 would be written as $v_j'(w_j - c_j - \psi_j) = \mathbf{E}v'(w_j - c_j + \varepsilon_j)$. Therefore, the results presented in this section, using the compensating precautionary premium, are equivalently obtained using the equivalent precautionary premium. Just the interpretation would change.

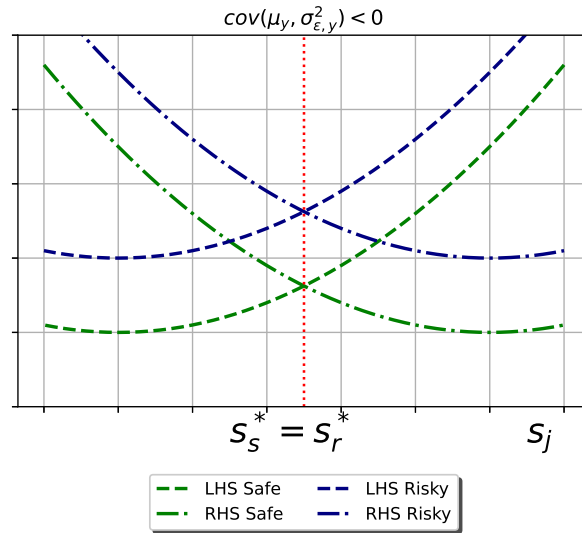


FIGURE 3. Euler equations with different income levels and risk.

risky consumer is slightly larger than what is assumed in Figure 3, such that she has a slightly higher amount of savings (i.e.; the risky RHS is higher)? The answer to this question would be found again in the negative correlation between income level and income risk. A higher level of income risk would correspond to a lower income level, bringing the consumption level of the risky consumer to an even lower level and therefore raising even more, her Euler Equation LHS. This would imply a further reduction of the amount of saving optimally chosen and eventually bringing it back to the same level of the safe consumer.

3.2. Prudence and the compensating precautionary premium

I am now in a position in which I can place this simple application within the (Kimball 1990b)'s model. To this end, I will use equation 11 of his paper:

Given the two second-period utility functions v_s and v_r if $\eta_r(s) = \frac{-v_r'''(s_r)}{v_r''(s_r)} > \eta_s(s) = \frac{-v_s'''(s_s)}{v_s''(s_s)}$ then, it must be the case that for all levels of saving and risk we have that:

$$(13) \quad \psi_r^*(s_r, \varepsilon_r) > \psi_s^*(s_s, \varepsilon_s)$$

Recall that $s_j = w_j - c_j$ and the risky consumer has a lower consumption level. Therefore she has a higher marginal utility compared to the safe consumer. This implies that the risky consumer is more absolutely prudent than the safe consumer, $\eta_r(s) > \eta_s(s)$, even if the two consumers have the same level of optimal savings as depicted in Figure 3.

Still, note that preferences are the same. Indeed, the coefficient of relative prudence (CRP) is the same between the two consumers. For instance, under CRRA, since the consumption levels cancel out, we would have $CRP = \gamma + 1$ for both, with γ measuring the relative risk aversion. Notice also that the decreasing absolute prudence implies that $\psi_j^*(s, \varepsilon_j)$ is decreasing in s ³⁵. To see this, assume that $\psi_j^*(s, \varepsilon_j)$ is differentiable. Then, differentiating equation 12 for s yields:

$$(14) \quad \frac{\partial \psi_j^*}{\partial s_j} = \frac{[(-v''''(s_j)v''(s_j)) + v'''(s_j)^2] \sigma_{\varepsilon,j}^2}{\underbrace{v''(s_j)^2}_{<0} 2}$$

Therefore, $\psi_j^*(s, \varepsilon_j)$ is decreasing in s .

3.3. The effect of income risk on the MPC

Given that $\psi_j^*(s, \varepsilon_j)$ is the amount of additional wealth needed to insure against risk, or the level of additional wealth to induce consumption in the presence of risk, we have that:

$$(15) \quad w_j(c, \varepsilon_j) = w_j(c_j, 0) + \psi_j^*(\varepsilon_{y,j}, \underbrace{w_j(c_j, 0) - c_j}_{\text{saving}})$$

If ψ_j^* is differentiable, then we can differentiate equation 15 and get the effect of income risk on the reciprocal of the MPC:

$$(16) \quad \frac{\partial w_j(c_j, \varepsilon_j)}{\partial c_j} = \frac{\partial w_j(c_j, 0)}{\partial c_j} + \frac{\partial \psi_j^*(\varepsilon_{y,j}, s_j)}{\partial s_j} \frac{\partial s_j(c_j, 0)}{\partial c_j}$$

Now, given the concavity of the consumption function and additive separability of the utility function across periods, we have that $\frac{\partial s_j(c_j, 0)}{\partial c_j} > 0$ ³⁶. Therefore the effect of income risk $\varepsilon_{y,j}$ on the reciprocal of the MPC $\frac{\partial w_j}{\partial c_j}$ is given by the sign of $\frac{\partial \psi_j^*}{\partial s_j}$. In contrast, the effect of income risk on the MPC is given by the opposite sign of $\frac{\partial \psi_j^*}{\partial s_j}$. Because absolute prudence is decreasing in s or alternatively in c , we have that $\frac{\partial \psi_j^*}{\partial s_j} < 0$. As a result, income risk increases the MPC with respect to the certainty equivalent.

However, notice that $\psi_r^* > \psi_s^*$ does not imply that risky households have higher MPC

³⁵See also (Kimball 1990a,b) for a more detailed discussion on this.

³⁶Alternatively, second-period consumption is a normal good.

than safe households.

Lemma 1 - The Risky consumer has a lower MPC: Assuming that absolute prudence is decreasing in s or c , given two consumers $j = s, r$, again, defined as $s =$ safe consumer, $r =$ risky consumer, having total wealth given by $w_j = w_0 + \mu_{y,j}$, non-capital income given by $y_j = \mu_{y,j} + \varepsilon_j$, where w_0 is the same for both and $\mu_{y,s} > \mu_{y,r}$, and $\sigma_{\varepsilon,s}^2 < \sigma_{\varepsilon,r}^2$; if proposition 1 and 2 hold true, that is: there exists some compensating precautionary premium such that $\psi_r^* > \psi_s^*$ and $s_r^* = s_s^*$ then, the risky consumer will have a lower MPC than the safe.

To see this, note that the MPC is given by:

$$(17) \quad \frac{\partial c_j}{\partial w_j} = \frac{1}{\frac{\partial w_j(c_j, \varepsilon_j)}{\partial c_j}} = \frac{1}{\frac{\partial w_j(c_j, 0)}{\partial c_j} + \frac{\partial \psi_j^*(\varepsilon_{y,j}, s_j)}{\partial s_j} \frac{\partial s_j(c_j, 0)}{\partial c_j}}$$

Note also that by the inverse function theorem, the effect of $\frac{1}{\frac{\partial \psi^*}{\partial s}}$ on the MPC is positive. It is now obvious that the denominator in equation 17 is larger for the risky consumer, lowering her MPC compared to the safe consumer. To see this, recall Equation 14. Alternatively, note that $\eta_r(s) > \eta_s(s)$ which implies $\psi_r^*(\varepsilon_r, s_r) > \psi_s^*(\varepsilon_s, s_s)$, and therefore $\frac{\partial \psi_r^*}{\partial s_r} > \frac{\partial \psi_s^*}{\partial s_s}$ holds true.

3.4. Wealth vs. Risk effects

So far, I established that given the negative correlation between income risk and income level, the low-income and high-risk consumer (risky) has a lower MPC than the safe consumer.

To shed some light on the various channels at play and further understand the mechanisms, let us ignore the negative correlation between income level and income risk. To this end, assume two consumers have the same expected income μ_y , and some different mean preserving spread over second-period consumption exists.

As shown in Figure 4 (left), the LHS of the Euler equation would be the same among the two consumers. However, the right-hand side would shift up for the consumer who is facing higher uncertainty. This is the well-known *Rothschild-Stiglitz effect* that arises when the comparison is made between two uncertain worlds, all else being equal. In this case, the entire problems move along the same LHS, and the consumer who is facing more

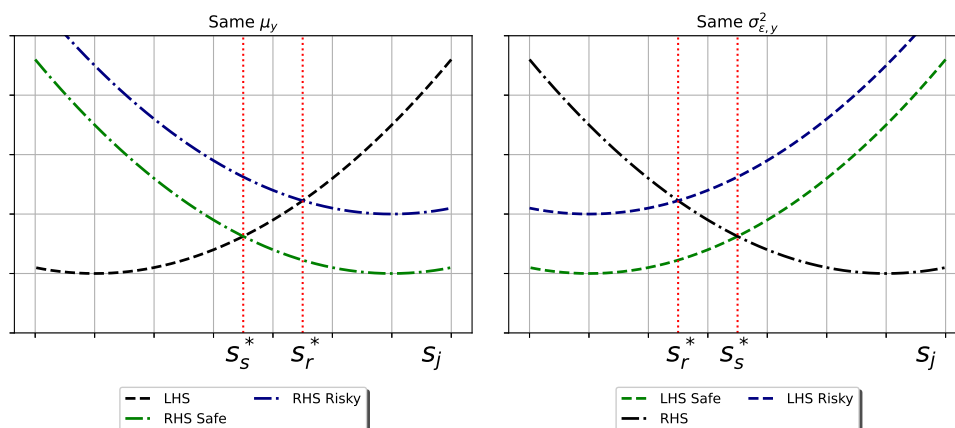


FIGURE 4. LEFT: Euler equations with different income risk only.
 RIGHT: Euler equations with different income levels only.

uncertainty would also be richer as she optimally chooses a higher level of s^* . In this case, given the concavity of the consumption function, the MPC of the new risky consumer would be lower mainly because of a wealth effect induced by the fact that now, the higher precautionary motive is not partially offset by a lower income level.

In contrast, suppose now that the two consumers have the same risk but different income levels. As shown in Figure 4 (right), in this case, there will be a clear income effect on the level of saving, as now the problems move along the same RHS of the Euler equation. The overall effect on the MPC is now ambiguous. Because the consumption function is concave, on the one hand, the MPC is higher at a lower wealth level. On the other hand, the MPC is lower at higher levels of prudence. To see this, note from equation 14 that even if the consumers share the same level of risk, a lower initial level of income implies a lower first-period consumption, implying a higher first-period marginal utility. This also implies that the low-income consumer has a higher coefficient of absolute prudence. Therefore, the denominator in equation 17 would still be higher for the low-income consumer. However, as shown in Figure 4, more is needed. Still, the two consumers would have different levels of wealth.

Finally, note that whenever the two consumers are not following their own Euler equation, i.e., $s_j^* = 0$, the MPC equals 1, regardless of income risk. This fact explains why the MPC differential among safe and risky households with low liquid assets is not statistically different from zero in the data.

3.5. Main Lesson

The negative correlation between income level and income risk creates heterogeneity in absolute prudence and homogeneity in savings levels between two consumers, ensuring a comparison within a specific point of wealth and facilitating the identification of the decreasing effect of income risk on the MPC. Therefore, it is important to clarify that the impact of income risk on the MPC does not solely depend on the level of income risk itself. Rather, the effect is identified thanks to the interaction between income level and income risk, which allows to shut down eventual wealth effects.

Additional Theoretical Results The online Appendix F also provides a proof of the results for an infinite-period model, exploiting the continuous time closed form solutions presented by (Achdou et al. 2021).

4. Quantitative analysis: A standard Incomplete Market model with ex-ante heterogeneity in the income processes

In this section I extend the intuition derived from the simple two-period model above to a an infinite-horizon Heterogeneous Agent Standard Incomplete Market Model (SIM), augmented with ex-ante heterogeneity in the income process.

Framework - The economy is populated by a continuum of households of measure one that is heterogeneous in their wealth a and income y . Time is discrete. Ex-ante heterogeneity in income risk is modeled by splitting the continuum into two permanent³⁷ groups: a share $(1/2)$ of *risky* households with higher income risk and a complementary share of *safe* households with lower income risk.

Both groups of household $j = \{S = Safe, R = Risky\}$ maximize their expected utility according to:

$$(18) \quad V(s_j, a_j) = \max_{c_j, a'_j} u(c_j) + \beta \mathbb{E}[V(s'_j, a'_j) | s_j]$$

$$(19) \quad \text{s.t. } a'_j + c_j = (1 + r)a_j + y_j(s_j)$$

$$(20) \quad a'_j \geq 0$$

³⁷I assume that the transition rate between income risk profiles is zero. I envisage whether and when households move from risky to safe or vice-versa in the data. Therefore, this assumption does not imply any loss of generality when comparing the MPCs from the model to those from the data.

(21) where:

$$(22) \quad u(c) = \begin{cases} \frac{c_j^{1-\gamma}}{1-\gamma} & \text{if } \gamma \neq 1 \\ \ln(c) & \text{if } \gamma = 1 \end{cases}$$

The income $y_j(s_j)$ is stochastic and depends on some discretized exogenous state s_j . To make the notation less cumbersome, I write y_j hereafter. The right-hand side of equation 19 is cash-on-hand, while the left-hand side is given by tomorrow's assets, a'_j , and current consumption c_j . To isolate the main mechanisms, I assume that households cannot borrow. However, departures from this assumption keep the results the same. A household of type j has a log income $y_{j,t}$ whose process is given by the sum of two orthogonal components, an AR(1) component, and an IID component.

$$(23) \quad \log(y_{j,t}) = z_{j,t} + \varepsilon_{j,t}$$

$$(24) \quad \begin{aligned} z_{j,t} &= \phi_j + \rho z_{j,t-1} + \eta_{j,t} \\ \mu_{z,j} &= \phi_j / (1 - \rho_{z,j}) \end{aligned}$$

Where $\varepsilon_{j,t}$ and $\eta_{j,t}$ are independently and identically distributed normal processes with mean zero and respective standard deviations of $\sigma_{j,\varepsilon}$ and $\sigma_{j,\eta}$. The term ϕ_j is a constant, with $\phi_S > \phi_R$. This assumption allows me to explicitly capture the negative correlation between income and income risk. Accordingly, insofar $\phi_S > \phi_R$, then from equation 24, $\mu_S > \mu_R$ ³⁸. Consistently with the data and estimates of the permanent and transitory components for both groups, discussed below, I set $\sigma_{R,\varepsilon} > \sigma_{S,\varepsilon}$ and $\sigma_{R,\eta} > \sigma_{S,\eta}$.

Let $c(a, z, \varepsilon)$ and $a(a, z, \varepsilon)$ be the optimal decision rules. Normalizing income to 1, a 10% increase in ε_t ($\bar{\varepsilon} = 0.1$) is similar to a standard tax rebate. Using the budget constraint, we have that:

$$(25) \quad MPC_j(a, z)^* = \frac{c(a, z_j, \bar{\varepsilon}) - c(a, z_j, 0)}{\exp(z_j)\bar{\varepsilon}}$$

The total size of the income grid for each group j is 7×3 . The income processes are discretized using the (Rouwenhorst 1995) method. The stationary distribution is found using the *lotteries* algorithm, following (Auclert et al. 2021).

Decomposing Income Risk - As discussed in Section 2.1, an important shortcoming

³⁸Results are quantitatively unaffected when I let $\mu_S > \mu_R$ driven by $\rho_S \geq \rho_R$.

of the CV is that it captures total income risk, hence, does not distinguish between permanent and transitory income shocks. However, for practical purposes, a useful exercise quantifies the variance of both components for risky and safe households.

To gather this information, I adopt a simple but coarse strategy. Essentially, I estimate the auto-covariance matrix of the stochastic process described in Equation 4 for both groups of households. The estimation strategy is standard and based on minimizing the distance between the empirical covariance matrix of income residuals and the theoretical counterpart implied by the income process. It should be clear that this exercise provides the variance of both components coming from realized income flows for households with different expected income flows. Therefore, these are not the variances of the expected components. Details are discussed in the Online Appendix C.

As expected, risky households display a higher variance in both (realized) components, especially in the permanent component. For both groups, the variance of the transitory component tends to be pro-cyclical. In contrast, the variance of the permanent component is quite counter-cyclical, reflecting the rise in risk inequality arising during recessions documented in the Online Appendix A. Risky households have an average biannual variance of 0.146 for the permanent shocks and 0.207 for the transitory shocks. Safe households have much lower average variances: 0.034 for the permanent shocks and 0.111 for the transitory shocks. Average population estimates align with (Jappelli and Pistaferri 2010a).

TABLE 8. Parameter Values: Biannual Calibration

Parameters	Value	Description	Traget
r	0.025	Interest rate	Biannual rate
γ	1	Elasticity of intertemporal substitution	
ρ_{zj}	0.986	Persistence of permanent shock	SHIW estimates
σ_{ps}	0.185	Standard deviation of Permanent shock, Safe hhs	SHIW estimates
σ_{pr}	0.383	Standard deviation of Permanent income, Risky hhs	SHIW estimates
σ_{ϵ_s}	0.334	Standard deviation of transitory income, Safe hhs	SHIW estimates
σ_{ϵ_r}	0.456	Standard deviation of transitory income, Risky hhs	SHIW estimates
ϕ_s	0.014	Permanent wage premium, Safe hh	$(\Delta\mu_z = 700Euro)$
ϕ_r	0.0042	Permanent wage premium, Risky hhs	$(\Delta\mu_z = 700Euro)$
a_{\min}	0.0	Minimum value of assets grid	
β	0.8586	Discount Factor	$(A_j/y_j = 0.57)$

The parameter ϕ_j is calibrated to match an average difference in monthly permanent

income(defined as in Table 7), $\Delta\mu_z$ of around 700 Euros³⁹.

Parametrization - Consistently with the data structure, the baseline calibration is biannual. Parameters are reported in Table 8. The discount factor, β , is internally calibrated to match both groups' average liquid assets over income. These values are recovered from Table 3, yielding 0.57 for both groups of households. Interestingly, the discount factor is forced to be the same between the groups so that heterogeneity in preferences does not drive differences in assets and MPC.

Results- The outcome for the baseline case is plotted in Figure 5 and reported in the first row of Table 9. As shown in Figure 5, the distribution of liquid assets is not polarized, supporting the intuition derived in Section 3. Moreover, as shown in Table 9, the model perfectly matches the average assets-to-income ratio and MPC for both groups of households. The MPC shown in Figure 5 confirms the evidence in the data and supports that for households with low liquid assets, there are no differences in MPC driven by Income risk.

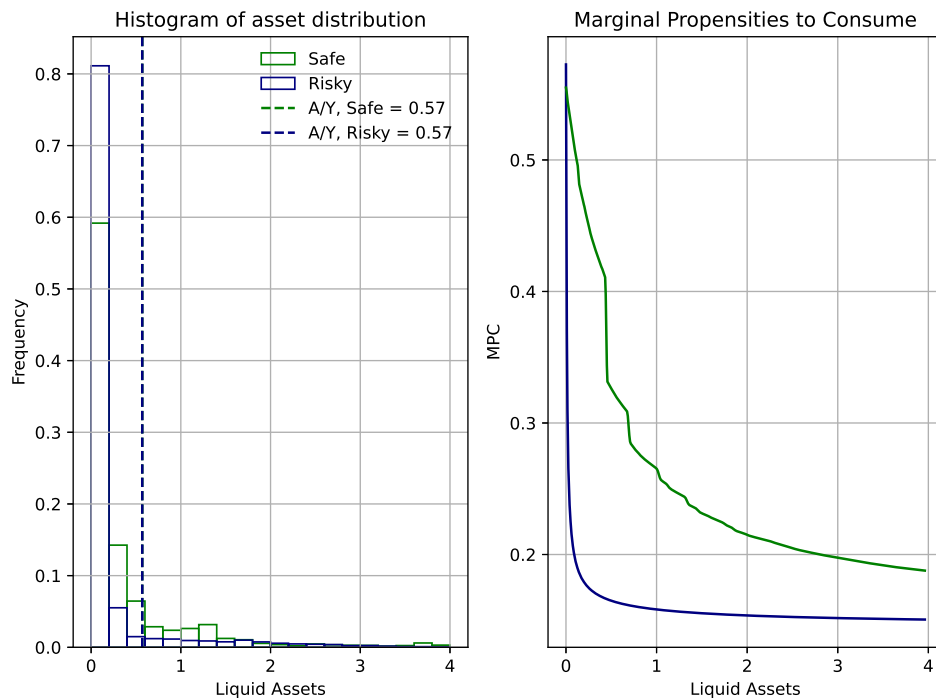


FIGURE 5. Summary Results baseline model.

Further, as shown in Table 9, I calibrate the model also at annual and quarterly frequencies. In both cases, assets are expressed as multiples of biannual ratios. Therefore,

³⁹The average permanent monthly income is 2.526,67 Euro for safe households, and 1.841,25 Euro for risky households.

both the interest rate and the discount factor are recalibrated accordingly. Variances of both components are adjusted accordingly. The annual variances are obtained by halving the biannual estimates. For quarterly variances, I perform two separate exercises. The first re-scales annualized variances following (Carroll et al. 2017). This formula divides by four the annualized variances of the permanent shocks and multiplies by four the variances of the transitory shocks. The persistence of the AR1 process is the same as biannual and set to 0.986.

TABLE 9. SUMMARY RESULTS, MODEL vs. DATA.

	Risky		Safe	
	Model	Data	Model	Data
Biannual MPC	0.329	0.289	0.473	0.503
Annual MPC	0.135	-	0.288	-
Quarterly MPC*	0.052	-	0.092	-
Quarterly MPC**	0.037	-	0.105	-
Liquid Ratio (A/Y)	0.57	0.57	0.57	0.57
Discount Factor (biannual)	0.8549		0.8549	
Discount Factor (annual)	0.9477		0.9477	
Discount Factor* (quarterly)	0.9878		0.9878	
Discount Factor** (quarterly)	0.9873		0.9873	
Income gap*** ($\Delta\mu_z$)				
Biannual	0.676			
Annual	0.829			
Quarterly*	0.737			
Quarterly**	0.842			

Data refers to the group-average MPC from Table 3

Wealth in both annual and quarterly calibrations is expressed as a multiple of biannual statistics obtained from Table 2.

* Variances re-scaled according to the (Carroll et al. 2017)'s formula.

** Variances re-scaled according to the (Krueger, Mitman, and Perri 2016)'s formula.

*** Income gap refers to the necessary ($\mu_{z, safe} - \mu_{z, risky}$) to match wealth in the data once variances are re-scaled to fit the model's frequency.

The second quarterly calibration uses the (Krueger, Mitman, and Perri 2016)'s formula. In this case, the formula for obtaining the quarterly variances of the tran-

sitory shocks is: $\sigma_{j,\varepsilon}^2 = (\sigma_{j,\varepsilon}^{\text{ann}})^2$. For permanent shocks, the transformation requires: $\sigma_{j,\eta}^2 = \left(\frac{1-\rho_{z,j}^2}{1-(\rho_{z,j}^{\text{ann}})^2}\right) (\sigma_{j,\eta}^{\text{ann}})^2$. Where $\rho_{z,j} = (\rho_{z,j}^{\text{ann}})^{0.25}$, is the transformed persistence of the AR1 process. Results are quantitatively similar in both cases and preserve the fact that risky households have a lower MPC.

Interestingly, as seen from Table 9, once variances are re-scaled to fit the new model's frequency, updating the new corresponding average income level to match the same wealth-to-income ratio is crucial. As discussed in Section 3, the negative correlation between income level and risk implies that if variances are reduced (increased), average permanent income must increase (decrease). This adjustment re-scales the Euler equations properly and prevents the possibility that the results are driven by ad-hoc calibrations of the magnitude of both components' variances. All calibrations confirm the empirical results and the theory presented in Section 3.

4.1. Robustness: Wealth, Risk, and Preferences effects

The theory presented in Section 3 suggests that whenever the negative correlation is neglected, it is impossible to compare the MPCs within-wealth. Therefore it is hard to mute the effects coming from the concavity of the consumption function. To check this prediction, I perform two additional exercises. The negative correlation between income risk and income level is neglected in both cases.

Table 10 shows the case in which $\Delta_{\mu,z,j} = 0$ That is, the wage premium ϕ_j is equal for both masses of households. All other parameters are as in the baseline case.

TABLE 10. SUMMARY RESULTS, BIENNIAL MODEL WITH SAME AVERAGE PERMANENT INCOME.

	Same Mean			
	Risky		Safe	
	β fixed	A/Y fixed	β fixed	A/Y fixed
Biannual MPC	0.335	0.566	0.473	0.473
Liquid Ratio (A/Y)	1.70	0.57	0.57	0.57
Discount Factor	0.8549	0.7201	0.8549	0.8549

To discipline the analysis and check whether this model version can match both liquid wealth and MPCs as in the data - and crucially given the same preferences -as in the baseline model, I require that the discount factor between the groups is the same.

Results for this case are shown in Table 10 under the column β fixed, namely, the same discount factor for both safe and risky households. The theory presented in Section 3 is

perfectly verified. Given the same preferences and same average income, households with higher income risk have much higher wealth. Notice that the liquid assets over income ratio rise to 1.70 for risky households. Interestingly, they now have almost the same MPC as in the baseline case. However, because now they are wealthier, it is impossible to isolate the risk effect. This result aligns with the buffer and stock theory, (Carroll 2001). In this case, households with higher income risk have lower MPC because their target wealth is higher. Further, notice that this is the same exercise done by (Zeldes 1989,a). If I set a smaller income risk difference between the two groups, such that risky households do not accumulate so much wealth, the MPC difference between safe and risky would be smaller. This fact will be clearly shown in Figure 7 below.

Therefore, the outcome of Table 10 suggests that this version of the model cannot match the wealth in the data given the same preferences. This fact implies that heterogeneity in discount factors is necessary to match the wealth in the data. The column A/Y fixed does this job. I require that both groups have the same liquid assets over income ratio as in the data. This is possible only if I let risky households be much more impatient. Their required biannual β now must be 0.7201. Their MPC, now, is driven by a preference effect. They are impatient and have a higher MPC compared to safe households.

Taking stock, for the case in which average permanent income equals and assuming only heterogeneity in income risk, the model matches the MPC or the wealth, but not both. This fact leads either to a wealth effect or to a preference effect.

TABLE 11. SUMMARY RESULTS, BIENNIAL MODEL WITH SAME RISK.

	Same Risk, $\Delta\mu_{z,j} = 0.676$			
	Low permanent Income		High permanent Income	
	β fixed	A/Y fixed	β fixed	A/Y fixed
Biannual MPC	0.473	0.282	0.473	0.473
Liquid Ratio (A/Y)	0.18	0.57	0.57	0.57
Discount Factor	0.8549	0.9216	0.8549	0.8549
	$\Delta\mu_{z,j}$ Adjusts			
Biannual MPC	0.473	-	0.473	-
Liquid Ratio (A/Y)	0.57	-	0.57	-
Discount Factor	0.8549	-	0.8549	-
Required $\Delta\mu_{z,j}$	0.00			

The second exercise, instead, assumes that heterogeneity in permanent income gives the only difference between the two masses. Risky households are now called "Low-permanent

income," and safe households are called "high-permanent income." As before, I set the permanent income difference between the two groups to match a gap of around 700 Euros. Results are summarized in Table 11. As before, I start requiring no preference heterogeneity; that is: β is the same for both masses of households.

If MPCs are driven by permanent income only, it must be the case that this version of the model can match the data and confirm the empirical evidence, rejecting the theory I illustrated in Section 3.

However, this is not the case. Again, given the same preferences, there is polarization in the stationary distribution of liquid assets. Consistently with the simple illustration of Section 3, households with low-permanent income have a lower amount of liquid wealth. The MPC is almost the same as households with higher permanent income (their problem is unchanged). Still, as shown in the column "A/Y fixed" of Table 11, I can reproduce the empirical results. However, this result is achievable only with preference heterogeneity.

Indeed, to accumulate the same wealth as in the data, the low-permanent income households must be much more patient. Their biannual discount factor must rise to 0.9216, lowering their MPC compared to high-permanent income households.

This fact proves that papers documenting that employed households with higher income (only) have higher MPCs, suffer from an omitted variable bias: the fact that high income negatively correlates with low risk.

To further illustrate this bias, I let the permanent income differential freely adjust to match the same liquid assets over income as in the data, given the same discount factor, i.e., meaning that unobserved heterogeneity is removed. The permanent income differential required to match the wealth data and remove the wealth polarization to allow for a within-wealth comparison of the MPC is zero. Moreover, the MPC between the groups is now equal. This means that once unobserved heterogeneity is removed (the β s are the same), differences in permanent income only have no effect on the MPC.

4.2. Generalization of the Results

To show that the results documented so far are general and not subject to a particular value of the income process parameters, I now perform some counterfactual exercises. Intuitively, by doing a simple comparative static analysis, I show the resulting MPCs and target wealth (defined as the liquid asset over income) achieved in the model's steady state for a wide range of income risk values. This type of analysis is important because it allows us to understand the dynamics behind the results documented so far.

Figure 6 displays the target wealth as a function of different values of the standard

deviation of permanent shocks⁴⁰ for both groups of households. The figure shows the three cases considered before: the baseline, in which a negative correlation exists between income risk and income level. The second case: in which income levels are the same. The third case: in which income risk is the same. To enhance the visualization of the results shown so far, Figure 6 also contains the outcome of the baseline parametrization of the model for each case as represented by green and blue dotted points.

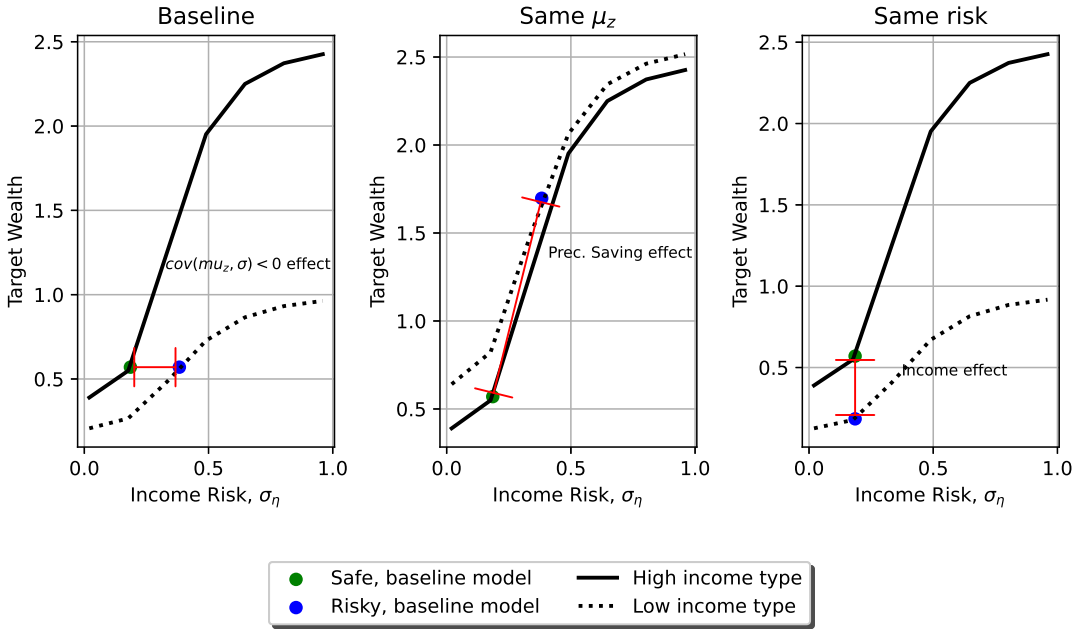


FIGURE 6. Target Wealth as a function of income risk: No preference heterogeneity, (β Fixed)

From the left panel, we can immediately see the implications of Section 3. The Euler equations heterogeneity is reflected in the asset accumulation dynamics for both groups. The asset risk curve for risky households is shifted to the right. This is the effect of the lower income level, which is why the two households have almost the same assets. Suppose now we assume that risky households face higher income risk than what is set in the baseline model and displayed by the blue dot. The income gap must rise to keep the negative correlation between income risk and income level constant, lowering and shifting the asset accumulation of risky households even further to the right. That is, a higher level of risk would correspond to a lower income level.

The wealth effect documented in the second model (same income) is reflected in the center panel of Figure 6. In this case, there are no offsetting forces given by the low-income levels. Thus the asset risk curve for the risky households is shifted up to the left. This is

⁴⁰Results are unchanged using the standard deviation of transitory shocks.

why risky households have much more wealth than safe households in the second case. This is a pure wealth (or precautionary wealth) effect.

Finally, in the third case, in which income risk is the same, we can see again the role played by a lower income level, which shifts the asset-risk curve to the right as in the baseline case. However, because the income risk level is the same for both groups of households, the low-income group ends up with lower assets. This is a pure income effect.

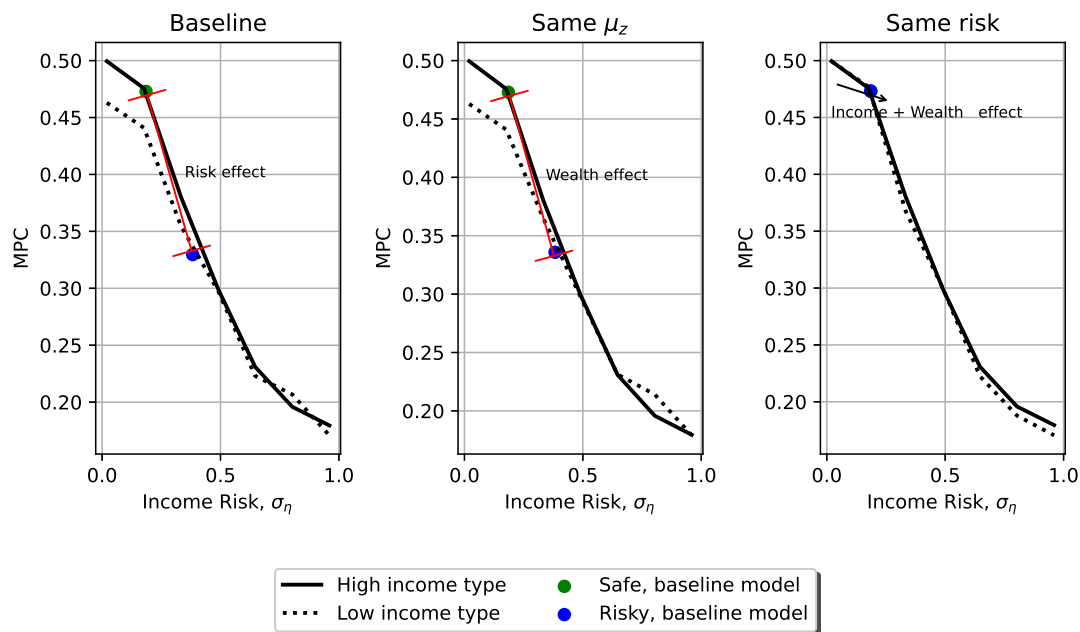


FIGURE 7. MPC as a function of income risk: No preference heterogeneity, (β Fixed)

Figure 7 repeats the same exercises plotting the MPC-risk curve and illustrating the dynamics at play for both groups in all three cases considered. In the baseline case, the effect of income risk is identified because wealth effects are muted. Notice that the MPC is decreasing in wealth, and wealth is increasing in risk. This is the endogeneity problem that masks the effects of income risk on the MPC, and that can be identified only by comparing MPC with wealth.

As predicted by the simple model of Section 3, Results are virtually the same in the center panel. Nevertheless, this is a wealth effect.

The right panel shows the case in which only income is different (same risk). This assumption creates an ambiguous effect on the MPC. On the one hand, low-income households have lower wealth. Thus their MPC is higher. On the other hand, the MPC of high-income households is also high for low levels of income risk (because their precautionary motive is weaker). Therefore the two groups have the same MPC.

4.3. One vs. Two-Asset models

One of the key features allowing the quantitative one-asset model presented in this section to generate high MPC compared to the data was that I calibrated the ratio of liquid assets over income. This approach, also known as *liquid calibration*, on the one hand, is consistent with the fact that liquid wealth is a better measure of the funds that households can readily use to smooth consumption; on the other hand, it implies that the model neglects the illiquid part of the household's wealth. The intuition behind this calibration is that targeting a lower aggregate wealth requires a low discount factor, which raises the MPC. As documented by (Kaplan and Violante 2021), all one asset models feature an *MPC-wealth* tension: if calibrated to match high MPCs as in the data, generate an understated wealth distribution, vice-versa, if calibrated to match total net-worth, generate low MPCs levels. This tension can usually be solved in two ways: The first is to assume some form of ex-ante heterogeneity (i.e., discount factor, relative risk aversion). The second is to use a two-asset model. However, the drawback of the first approach is that the wealth distribution is polarized. For example, assuming ex-ante heterogeneity in discount factors implies that more patient households would hold all the wealth. Therefore, the model would miss the middle of the distribution.

Given the purposes of this paper, this fact is not important for documenting the results, mainly because of two reasons: First, as broadly documented so far, assuming ex-ante heterogeneity in the income processes and modeling the negative correlation between income level and income risk prevents the polarization of the wealth distribution. Second, despite targeting a higher amount of wealth, the ranking of the MPCs between the two household groups is still preserved. To see this point, note that when I calibrate the annual or quarterly version of the model, the liquid asset-to-income ratio is expressed as a multiple of biannual values. For the quarterly calibration, the stationary liquid wealth of safe and risky households goes from 0.57 to 4.56. Still, safe households had larger MPC than risky ones.

This fact means that despite the model can generate a large fraction of unconstrained households with large MPC, as with all one-asset models, it suffers the typical trade-off. This also requires the results to be validated through a two-asset model's lens. The Online Appendix G.4 shows this exercise and demonstrates that a two-asset version of the model presented in this section can document the empirical facts while matching the exact amount of median net worth (net-liquid plus net-illiquid holdings) for both groups of households.

Additional Numerical Results Another interesting result is given by the fact that to generate the results of this paper one does not need to assume a complicated income

process as I did here. Assuming a simple AR1 with positive mean $\mu_{z,j}$ and letting differences in it be driven by the persistence parameters, one yields the same results: Safe households have higher MPC, and the dynamics are exactly the same as the one documented for each case.

To isolate the main dynamics, the numerical analysis conducted so far has abstracted from the presence of borrowing constraints. In practice, households can borrow, and in particular, the SHIW reports that only a meager share of households, around 4%, in Italy is constrained, (Jappelli and Pistaferri 2020). Moreover, (Mistrulli et al. 2021), using a unique dataset from one of the largest banks operating in Italy, shows that households with higher income risk tend to be more constrained than households with safer income paths (i.e., permanent workers). To account for this evidence, Table A.18, of the Online Appendix G.2, shows that the numerical findings are robust even accounting for the presence of borrowing constraints and different degrees of heterogeneity in access to credit.

5. Fiscal Implications

This section aims to quantify the fiscal implication of the evidence documented so far. It should be noted that a policymaker barely or never can directly observe the income risk households face. Moreover, as pointed out by (Jappelli and Pistaferri 2020), unlike labor income, the level of liquid assets, or cash on hand, is hardly accurately observed by the government. Therefore, using the coefficient estimates from Tables 4 and 5 would not be realistically informative.

Modifying the exercise in Table 6, such that the estimate of the MPC for labor categories are liquid-wealth unconditional, I can design, instead, a credible experiment. Indeed, a policymaker can easily observe the net labor income and the labor market statuses of workers. To this end, I simulate a policy that transfers the equivalent of the 1% of national income (in equal amounts) to the bottom x% of the current income distribution and assumes that specific labor categories receive the transfer. This experiment is very similar to various fiscal rebates provided by governments to cope with the financial aftermaths of the COVID-19 Pandemic. Consistent with this scenario, the baseline experiment assumes that the government targets the median income distribution of Self-Employed households, the most affected labor category.

Results are shown in Table 12. The aggregate response decreases in the percentile of labor income for all groups of households. However, the intercept is quite different across groups. The baseline scenario, targeting the median income distribution of self-employed

households, yields an aggregate demand increase of 0.15.

Consistently with the MPC estimates, this is the lowest aggregate demand (AD) response. What would be the gain of giving another category the same amount of money?

TABLE 12. AGGREGATE (p.a.) RESPONSES OUT OF A RE-DISTRIBUTIVE POLICY

Recipients	All categories	Civil Servants	Private sector	Self-Employed
Target, $y_i <$:				
P10	0.34	0.44	0.31	0.21
P25	0.31	0.40	0.28	0.19
P50	0.25	0.32	0.23	0.15
P75	0.20	0.25	0.18	0.12
P90	0.17	0.22	0.15	0.10
Random < P90	0.17	0.22	0.15	0.10
Target: Median Income Self-Employed				0.15
Counterfactuals	All categories	Civil Servants	Private sector	
$\overline{\Delta AD}$	+66%	+113%	+53%	

H-Sample 1998-2016. Column 2 uses the estimates for the risk unconditional MPC, while columns 3-5 use the wealth unconditional MPC, shown in the Online Appendix B.9.

Aggregate demand responses are weighted by the relative share of labor categories over the total employed workforce.

The counterfactual exercises show that targeting the median of the labor income distribution only, without selecting a specific category, would yield a stronger AD effect, +66%⁴¹, compared to the effect obtained by targeting self-employed households. The "*paradox of risk*" can be quantified if (absurd) the government gave the same amount of money to the safest category, the Civil Servants. In that case, the AD effect would be +113% compared to the baseline scenario.

Of course, the message of this back-of-the-envelope exercise is not to say that the government's decision should have been unpopular. Rather, this simple example provides an approximate quantification of the costs associated with income risk, suggesting that a government cannot simultaneously, or using the same fiscal tool, boost aggregate demand and mitigate income risk.

The online Appendix G.3 repeats this exercise through the lens of the model presented

⁴¹ $(0.25/0.15)-1=0.66$.

in section 4 but extends it to account for general equilibrium feedback and compute the multipliers of this policy. Results confirm the intuition provided in this simple exercise.

6. Conclusion

This paper has provided empirical, theoretical, and numerical evidence to understand the effects of labor income risk heterogeneity on the MPCs. It has been shown that when income levels and income risk negatively correlate, a higher level of income risk lowers the MPC. This effect is identified thanks to the interaction between income level and income risk, which allows to shut down eventual wealth effects. It has also been shown that neglecting this interaction, wealth effects drive the MPC, masking the effects of risk. Further, it has been shown that the MPC does not depend on (permanent, persistent, or current) income levels themselves, whose direct effect on the MPC is always ambiguous.

The overall evidence suggests a significant amount of heterogeneity within the wealth distribution. In particular, the Italian economy has a considerable fraction (25-38%) of unconstrained households facing low income risk and exhibiting higher biannual MPCs (out of transitory shocks) than households facing higher levels of income risk. This fact has been rationalized through the lens of standard models with precautionary saving.

Finally, the evidence found in this paper highlights challenges for the design of targeted stimulus checks aimed at restoring particular labor categories facing substantial labor uncertainty. In these cases, a government faces a trade-off between mitigating income risk and maximizing the aggregate demand response of the fiscal expansion.

Future studies may explore this trade-off by comparing stimulus checks with unemployment insurance extensions. At the same, extending the model presented in Section 4 to study the monetary transmission to individual consumption is quite straightforward. However, this analysis requires deeper details, therefore is left for future research.

Acknowledgments. This version supersedes an earlier version of the paper entitled “Within Wealth MPC Heterogeneity” It is supplemented by an online appendix <https://www.ettoresaveoia.me/research/working-papers>. Replications materials are made available on <https://www.ettoresaveoia.me/research/working-papers>

This paper has been concluded during my research period at the Directorate General Research of the European Central Bank and has been part of the work on the Research Task Force on Heterogeneity in Macroeconomics.

I am indebted to Tullio Jappelli and Alberto Zazzaro for their invaluable support. This paper particularly benefited from discussions and comments provided by Christopher Carroll, Bence Bardoczy, Francesco Caruso, Oreste Tristani, Micol De Vera, Riccardo Cioffi, Gianluca Violante, Ben Moll, Jirka Slacalek, Jeanne Commault, Moritz Lenel, Ralph Lutticke.

The views expressed in this paper are those of the author and do not necessarily reflect the views of the European Central Bank, the ESCB, or the University of Naples - Federico II.

Author Ettore Savoia: Department of Economics and Statistics - University of Naples Federico, Via Vicinale Cupa Cintia, 26, 80126 Napoli (NA), Italy ettore.savoia@unina.it. Any remaining errors are my own.

References

- Abowd, John M., and David Card. 1989. “On the Covariance Structure of Earnings and Hours Changes.” *Econometrica* 57 (2): 411–445.
- Achdou, Yves, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. 2021. “Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach.” *The Review of Economic Studies* 89 (1): 45–86.
- Aguiar, Mark A, Mark Bils, and Corina Boar. 2020. “Who Are the Hand-to-Mouth?” *National Bureau of Economic Research, Inc* (26643).
- Andreolli, Michele, and Paolo Surico. 2021. “Less Is More: Consumer Spending and The Size Of Economic Stimulus Payments.” *CEPR* (Dp15918).
- Arellano, Manuel, Stéphane Bonhomme, Micole De Vera, Laura Hospido, and Siqi Wei. 2022. “Income Risk Inequality: Evidence From Spanish Administrative Records.” *Quantitative Economics* 13 (4): 1747–1801.
- Auclert, Adrien, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub. 2021. “Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models.” *Econometrica* 89 (5): 2375–2408.
- Bayer, Christian, Ralph Luetticke, Lien Pham-Dao, and Volker Tjaden. 2019. “PRECAUTIONARY SAVINGS, ILLIQUID ASSETS, and THE AGGREGATE CONSEQUENCES OF SHOCKS TO HOUSEHOLD INCOME RISK.” *Econometrica* 87 (1): 255–290.
- Berger-Thomson, L., Elaine Chung, and Rebecca Mckibbin. 2009. “Estimating Marginal Propensities To Consume In Australia Using Micro Data.” *Reserve Bank Of Australia*.
- Bertaut, Carol C., and Michael Haliassos. 1997. “Precautionary portfolio behavior from a life-cycle perspective.” *Journal of Economic Dynamics and Control* 21 (8-9): 1511–1542.
- Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2008. “Consumption Inequality and Partial Insurance.” *American Economic Review* 98 (5): 1887–1921.
- Boutros, Michael. 2021. “Bounded Intertemporal Rationality and The Marginal Propensity to Consume.” *Duke University*.
- Bunn, Philip, Jeanne Le Roux, Kate Reinold, and Paolo Surico. 2018. “The consumption response to positive and negative income shocks.” *Journal of Monetary Economics* 96 (C): 1–15.
- Carroll, Christopher. 2001. “A Theory of the Consumption Function, with and without Liquidity Constraints.” *Journal of Economic Perspectives* 15 (3): 23–45.
- Carroll, Christopher. 2009. “Precautionary saving and the marginal propensity to consume out of permanent income.” *Journal of Monetary Economics* 56 (6): 780–790.
- Carroll, Christopher, Martin Holm, and Miles Kimball. 2021. “Liquidity constraints and precautionary saving.” *Journal of Economic Theory* 195: 105–276.
- Carroll, Christopher, and Miles Kimball. 1996. “On the Concavity of the Consumption Function.” *Econometrica* 64 (4): 981–992.
- Carroll, Christopher, Jiri Slacalek, Kiichi Tokuoka, and Matthew White. 2017. “The distribution of wealth and the marginal propensity to consume.” *Quantitative Economics* 8 (3): 977–1020.
- Christelis, Dimitris, Dimitris Georgarakos, Tullio Jappelli, Luigi Pistaferri, and Maarten van Rooij. 2019. “Asymmetric Consumption Effects of Transitory Income Shocks.” *Economic*

- Journal* 129 (622): 2322–2341.
- Commault, Jeanne. 2022a. “How Do Persistent Earnings Affect the Response of Consumption to Transitory Shocks” *Sciences Po Discussion Paper version*.
- Commault, Jeanne. 2022b. “Does Consumption Respond to Transitory Shocks? Reconciling Natural Experiments and Semistructural Methods.” *American Economic Journal: Macroeconomics* 14 (2): 96–122.
- Deaton, Angus. 1991. “Saving and Liquidity Constraints.” *Econometrica* 59 (5): 1221–48.
- Fagereng, Andreas, Luigi Guiso, and Luigi Pistaferri. 2016. “Portfolio Choices, Firm Shocks and Uninsurable Wage Risk.” *Nber Working Paper* (22883).
- Fagereng, Andreas, Martin Holm, and Gisle Natvik. 2018. “MPC Heterogeneity and Household Balance Sheets.” *CEsifo Working Paper Series* (7134).
- Fisher, Jonathan, David S. Johnson, Timothy M. Smeeding, and Jeffrey Thompson. 2020. “Estimating the marginal propensity to consume using the distributions of income, consumption, and wealth.” *Journal of Macroeconomics* 65 (C).
- Friedman, Milton. 1957. “Introduction to "A Theory of the Consumption Function".”
- Fuster, Andreas, Greg Kaplan, and Basit Zafar. 2020. “What Would You Do with \$500? Spending Responses to Gains, Losses, News, and Loans.” *The Review of Economic Studies* 88 (4): 1760–1795.
- Ganong, Peter, Damon Jones, Pascal J Noel, Fiona E Greig, Diana Farrell, and Chris Wheat. 2020. “Wealth, Race, and Consumption Smoothing of Typical Income Shocks.” Working Paper 27552, National Bureau of Economic Research.
- Gottschalk, Peter, and Robert Moffitt. 2009. “The Rising Instability of U.S. Earnings.” *Journal of Economic Perspectives* 23 (4): 3–24.
- Guiso, Luigi, Tullio Jappelli, and Daniele Terlizzese. 1992. “Earnings uncertainty and precautionary saving.” *Journal of Monetary Economics* 30 (2): 307–337.
- Guiso, Luigi, Tullio Jappelli, and Daniele Terlizzese. 1996. “Income Risk, Borrowing Constraints, and Portfolio Choice.” *American Economic Review* 86 (1): 158–72.
- Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song. 2021. “What Do Data on Millions of U.S. Workers Reveal about Lifecycle Earnings Dynamics?” *Econometrica* 89 (5): 2303–2339.
- Guvenen, Fatih, Serdar Ozkan, and Jae Song. 2014. “The Nature of Countercyclical Income Risk.” *Journal of Political Economy* 122 (3): 621–660.
- Haider, Steven. 2001. “Earnings Instability and Earnings Inequality of Males in the United States: 1967-1991.” *Journal of Labor Economics* 19 (4): 799–836.
- Heaton, John, and Deborah Lucas. 2000. “Portfolio Choice and Asset Prices: The Importance of Entrepreneurial Risk.” *The Journal of Finance* 55 (3): 1163–1198.
- Hoffmann, Eran, Davide Malacrino, and Luigi Pistaferri. 2021. “Labor Market Reforms and Earnings Dynamics: The Italian Case.” *Imf Working*.
- Jappelli, Tullio, and Luigi Pistaferri. 2010a. “Does consumption inequality track income inequality in Italy?” *Review of Economic Dynamics* 13: 133–153.
- Jappelli, Tullio, and Luigi Pistaferri. 2010b. “The Consumption Response to Income Changes.” *Annual Review of Economics* 2 (1): 479–506.
- Jappelli, Tullio, and Luigi Pistaferri. 2014. “Fiscal Policy and MPC Heterogeneity.” *American*

- Economic Journal: Macroeconomics* 6 (4): 107–36.
- Jappelli, Tullio, and Luigi Pistaferri. 2020. “Reported MPC and Unobserved Heterogeneity.” *American Economic Journal: Economic Policy* 12 (4): 275–97.
- Kaplan, Greg, Benjamin Moll, and Giovanni Luca Violante. 2018. “Monetary Policy According to HANK.” *American Economic Review* 108 (3): 697–743.
- Kaplan, Greg, and Giovanni Luca Violante. 2010. “How Much Consumption Insurance beyond Self-Insurance?” *American Economic Journal: Macroeconomics* 2 (4): 53–87.
- Kaplan, Greg, and Giovanni Luca Violante. 2021. “Propensity To Consume In Heterogeneous Agent Models.” *Annual Reviews Of Economics (Forthcoming)*.
- Kaplan, Greg, Giovanni Luca Violante, and Justin Weidner. 2014. “The Wealthy Hand-to-Mouth.” *Brookings Papers on Economic Activity* 45 (1 (Spring)): 77–153.
- Kimball, Miles. 1990a. “Precautionary Saving and the Marginal Propensity to Consume.” (3403).
- Kimball, Miles. 1990b. “Precautionary Saving in the Small and in the Large.” *Econometrica* 58 (1): 53–73.
- Koo, Hyeng Keun. 1998. “Consumption and Portfolio Selection with Labor Income: A Continuous Time Approach.” *Mathematical Finance* 8 (1): 49–65.
- Kosar, Gizem, Davide Melcangi, Laura Pilossoph, and Laura Wiczor. 2023. “Stimulus and Insurance: The Marginal Propensity to Discharge Debt.” *Working Paper*.
- Krueger, Dirk., Kurt. Mitman, and Fabrizio. Perri. 2016. “Macroeconomics and Household Heterogeneity, in Handbook of Macroeconomics.” *Elsevier*.
- Kueng, Lorenz. 2018. “Excess Sensitivity of High-Income Consumers.” *The Quarterly Journal of Economics* 133 (4): 1693–1751.
- Lagrosa, Ivan. 2022. “Income Dynamics in Dual Labor Markets.” *Available at SSRN: <https://ssrn.com/abstract=4232078> or <http://dx.doi.org/10.2139/ssrn.4232078>*.
- Luengo-Prado, Maria José, and Bent E. Sørensen. 2004. “The Buffer-Stock Model and the aggregate propensity to Consume: A panel-data study of the US STATES.” *CEPR - Discussion Paper No. 4474*.
- MaCurdy, Thomas E. 1982. “The use of time series processes to model the error structure of earnings in a longitudinal data analysis.” *Journal of Econometrics* 18 (1): 83–114.
- Mian, Atif, Kamalesh Rao, and Amir Sufi. 2013. “Household Balance Sheets, Consumption, and the Economic Slump*.” *The Quarterly Journal of Economics* 128 (4): 1687–1726.
- Misra, Kanishka, and Paolo Surico. 2014. “Consumption, Income Changes, and Heterogeneity: Evidence from Two Fiscal Stimulus Programs.” *American Economic Journal: Macroeconomics* 6 (4): 84–106.
- Mistrulli, Paolo Emilio, Tommaso Oliviero, Zeno Rotondi, and Alberto Zazzaro. 2021. “Job Protection and Mortgage Conditions: Evidence From Italian Administrative Data.” *Workings Papers* (39).
- Pratt, John W, and Richard Zeckhauser. 1987. “Proper Risk Aversion.” *Econometrica* 55 (1): 143–54.
- Romer, Christina. 2021. “The Fiscal Policy Response To The Pandemic.” *Brookings Papers On Economic Activity*.
- Rouwenhorst, K. G. 1995. “Asset Pricing Implications Of Equilibrium Business Cycle Models.”

Princeton University Press: 294–330.

Sala, Hector, and Pedro Trivín. 2021. “Reported Mpc In The Presence Of Debt.” *Economics Letters, Elsevier* 205(C).

Shore, Stephen H., and Todd Sinai. 2010. “Commitment, Risk, and Consumption: Do Birds of a Feather Have Bigger Nests?” *The Review of Economics and Statistics* 92 (2): 408–424.

Zeldes, Stephen. 1989,a. “Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence.” *The Quarterly Journal of Economics* 104 (2): 275–298.

Acknowledgements

I am indebted to Tullio Jappelli and Alberto Zazzaro for their invaluable support. This paper particularly benefited from discussions and comments provided by Christopher Carroll, Bence Bardoczy, Francesco Caruso, Oreste Tristani, Gonzalo Paz Pardo, Micol De Vera, Riccardo Cioffi, Gianluca Violante, Ben Moll, Jirka Slacalek, Jeanne Commault, Moritz Lenel, Ralph Lutticke.

This version supersedes an earlier version of the paper entitled “Within Wealth MPC Heterogeneity”. It is supplemented by an online appendix <https://www.ettosavoia.me/research/working-papers>.

This paper has been concluded during my research period at the Directorate General Research of the European Central Bank and has been part of the work on the Research Task Force on Heterogeneity in Macroeconomics.

Ettore Savoia

Sveriges Riksbank, Stockholm, Sweden; email: ettore.savoia@riksbank.se

© European Central Bank, 2023

Postal address 60640 Frankfurt am Main, Germany

Telephone +49 69 1344 0

Website www.ecb.europa.eu

All rights reserved. Any reproduction, publication and reprint in the form of a different publication, whether printed or produced electronically, in whole or in part, is permitted only with the explicit written authorisation of the ECB or the authors.

This paper can be downloaded without charge from www.ecb.europa.eu, from the [Social Science Research Network electronic library](#) or from [RePEc: Research Papers in Economics](#). Information on all of the papers published in the ECB Working Paper Series can be found on the [ECB's website](#).

PDF

ISBN 978-92-899-6243-8

ISSN 1725-2806

doi:10.2866/07882

QB-AR-23-103-EN-N