The Earnings Elasticities to Aggregate Fluctuations in the Euro Area

Atanas Pekanov

January 2024

This paper documents how the earnings of different household groups vary with the change in aggregate GDP in the Euro area. These household earnings elasticities to aggregate economic fluctuations are a crucial moment for models with household heterogeneity – macroeconomics dynamics are transmitted differently throughout the economy when different household types react differently to macroeconomic shocks. This paper is the first to calculate the relative earnings elasticity of different income groups in the countries of the Euro area using three separate datasets and comparing the results. The three datasets deliver a consistent picture of higher earnings elasticities for the bottom income groups which decline along the income distribution. Importantly, the slope of these household betas – the change of individual earnings to aggregate earnings changes, is similar across two very different types of datasets – aggregated data on income groups from the WID and GRID databases, and microdata from the ECB HFCS dataset. This paper is the first to use the HFCS panel dataset to track earnings dynamics of households in the Euro area. It also decomposes the difference between the two types of estimations and shows the importance of individual level data for properly accounting for income mobility. I also analyse the correlation between earnings elasticities of different income groups and their MPC and show that the groups with the highest earnings elasticities γ on average also have higher MPCs. I show this leads to an amplifying multiplier effect which increases the baseline MPC in the economy by between 17% and 34% and therefore the aggregate multiplier by 6.7% to 13.5%.

Keywords: Income inequality, Income risk, Earnings elasticity, Workers betas, MPC

JEL: D31, E12, D31, D12

Note: This is a work-in progress version of this paper. Comments are welcome.

The author would like to thank for helpful discussions and comments to Georgi Kocharkov, Jirka Slacalek, Adrien Auclert, Jesus Crespo Cuaresma and Pirmin Fessler. I am grateful to the ECB HFCS Network for providing me the HFCS data. The views and opinions expressed in this paper are strictly those of the author and do not necessarily reflect the views of the Austrian Institute of Economic Research, the European Central Bank or the ESCB. The author takes responsibility for any errors.

I. Introduction

How does accounting for households' heterogeneity – and in particular inequality in income and wealth – change our approach to macroeconomics? This has been one of the most active fields of study for macroeconomics in recent years[1]. How does the business cycle affect the relative income of different income groups and what does that imply for the amplification of macroeconomic dynamics after shocks and specific policies? This paper revolves around this specific question – how much do different income groups gain or lose in times of economic expansion or contraction? It is a question about the relative earnings elasticity of different income groups to aggregate output fluctuations and is critical for understanding macroeconomic dynamics in models with household inequality and heterogeneity.

Why does the relative earnings elasticity of households' matter? The risk of income shocks has implications for the transmission of macroeconomic shocks and economic policies through a number of channels. Most importantly, cyclical income risks can give rise to precautionary savings, thereby amplifying the response of consumption to aggregate output changes. This income risk varies with the business cycle and the variation of this income risk leads to large drops in consumption during recessions, implying a much higher aggregate consumption volatility than in standard models (McKay, 2017). In a model with liquid and illiquid assets, households decide to hoard liquid assets because of this heightened uncertainty, which results not only in reduced consumption demand, but also in less illiquid physical investments (Bayer et al., 2015). The heightened uncertainty channel works through unemployment risk, which induces wealth-poor households to save even more, thereby driving aggregate demand lower and leading to a self-fulfilling fluctuation of unemployment (Heathcote and Perri, 2018).

In a wide class of new Heterogeneous Agents New Keynesian models (HANK), household heterogeneity amplifies both macroeconomic shocks and the effects of macroeconomic policies through the heterogeneity of household MPCs, in comparison to traditional models with only a representative household. Besides the differences in household MPCs however, the amplification of aggregate shocks through aggregate demand with household heterogeneity crucially depends on the relative sensitivity of the income of different household groups to output fluctuations.

As shown by Bilbiie (2018) and Werning (2015), the amplification would be strengthened if households with lower income and with lower liquid asset holdings, which typically have higher MPCs, are more affected in terms of labour market outcomes and wages to aggregate output changes – that is, if they e. g. lose more frequently their jobs during recessions or suffer disproportionally income losses. This is a question stated clearly going back to Keynes, who formulated it as follows: *"The amount that the community spends on consumption obviously depends on [...] the principles on which income is divided between the individuals composing it which may suffer modification as output is increased."*, Keynes (1936). In technical terms, Bilbiie (2008) shows the result that e. g. monetary policy shocks are amplified with agent heterogeneity only when the elasticity of income of the constrained high-MPC agent is above 1.

The measurement in this paper has important implications for the literature in macro, especially so for labor and business cycle analysis. I calculate the relative earnings elasticity of different income

[1] For a summary on microeconomic heterogeneity and its macroeconomic implications see Kaplan and Violante (2018). The repercussions of household heterogeneity on monetary policy transmission are explored e. g. in Kaplan et al. (2018), while the repercussions for fiscal policy are explored e. g. in Hagedorn et al. (2019). There are numerous other channels through which inequality can affect monetary and fiscal policy – Bartscher et al. (2022) explore how racial inequality is affected by monetary policy.

groups γ_i in the countries of the Euro area using three separate datasets. Figure 1 summarizes the main results out of my three estimations. It reports the earnings elasticity by income deciles or at specific income percentiles to aggregate real GDP changes – how much does the income of this income group react to changes in aggregate economic activity. All three sets deliver a similar pattern – the incomes of the lowest income groups are much more sensitive to aggregate output changes than those in the middle of the income distribution. There is a clear downward pattern of the sensitivity of individual income to aggregate output changes. These estimates are often called in the literature workers betas, but I follow the naming convention by Auclert (2019) and define them as γ_i. These patterns of decreasing γ_i along the income distribution hold up until the very high income bins, where the sensitivity again increases for the highest earning groups in two of my datasets.

These are the first results of this type for the Euro area countries, comparing both aggregated data by income groups and income changes data I compile from the panel element of a household survey. The results from the two types of sources are similar and also in line with previous estimates for the US, reported by Guvenen et. al (2017). The results estimated with the HFCS panel data point to higher earnings elasticities for the bottom income deciles than those from administrative data, as the panel data is suited to also take into account income mobility across deciles – it enables me to measure the income change for households which move along the income ladder, unlike data reported in aggregated form from administrative tax data as in the case of the WID and GRID database. The difference in the magnitude between the two approaches is due to this difference in measurement, but importantly, the downward sloping pattern persists across all three datasets. I confirm these results by running a number of robustness checks using different measurements for income and different sample selections.

To estimate how does this different income elasticity of individual earnings affect the aggregate economy, I then explore the correlation between earnings elasticities of different income groups and their MPC. My results show that the groups with the highest earnings elasticities on average also have higher MPCs – a finding also reported in Patterson (2022) for the US. In my baseline estimation, the covariance between earning elasticities of different income groups and their MPCs is around 6%. This finding has important repercussions for the reaction of consumption to macroeconomic shocks or different economic policies – the fact that individuals with the highest sensitivity to economic changes also have the highest MPCs act as an amplification mechanism to any initial shock or policy change. To evaluate this, I calculate the size of the aggregate MPCs with and without this amplification mechanism, following Patterson (2022). For my baseline specification, I find that the amplification mechanism increases the aggregate MPC of the economy by around 15%.

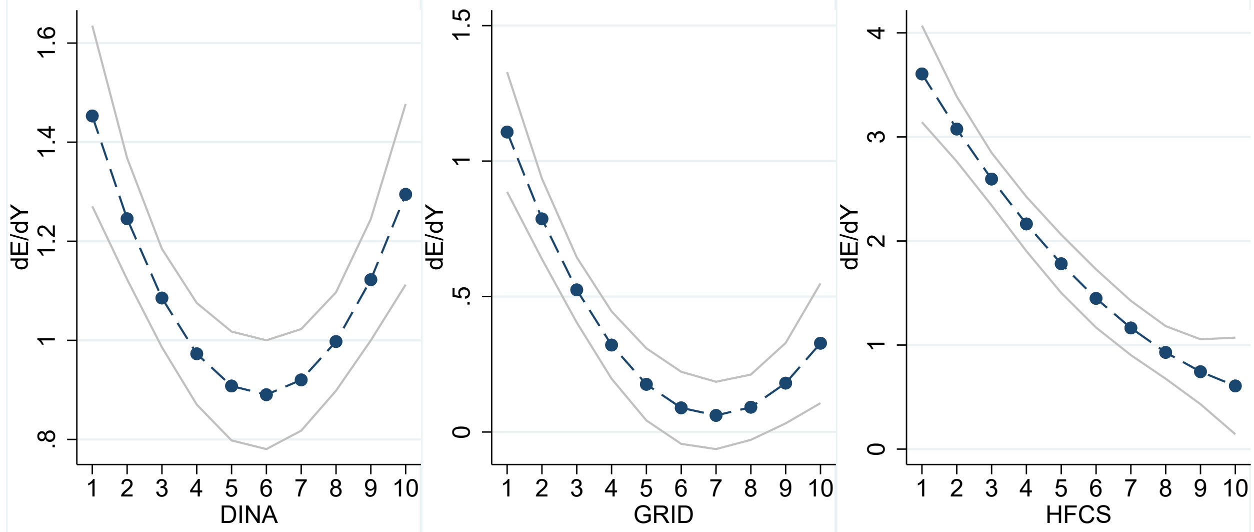

Figure 1: Earnings elasticity dE/dY by income decile estimated with three different datasets – Distributional National Accounts Data from the World Inequality Database (DINA), the Global Repository on Income Dynamics (GRID) and the ECB Household Finance and Consumption Survey (HFCS).
Note: The three graphs show the polynomial function based on ten estimations of the earnings elasticities of households to GDP changes from three different datasets. For the DINA (WID) dataset earnings elasticities are estimated at ten specific deciles. For the GRID dataset earnings elasticities are estimated at ten specific percentiles across the income distribution. For my main dataset – the HFCS earnings elasticities are estimated for the bin of each of the ten income deciles. The HFCS dataset uses panel data and I calculate the elasticities based on individual household dynamics, while the WID and GRID datasets report incomes at specific points of the distribution. Since the HFCS dataset has the additional advantage to permit for full income mobility across bins, the magnitudes are higher, whereas in the two other datasets calculations are always smoothed by mean reversion.

II. Literature

Macroeconomic aggregates are the sum of microeconomic variables describing the economic performance of agents in the economy, but changes in these macroeconomic aggregates may influence different agents to a varying degree. In the case of economic growth, this is normally transmitted throughout the economy to individuals and households by increasing their individual income. The standard New Keynesian model using a representative agent then requires that the aggregate income change translates into an income change for the sum of representative agents in this economy. This change is proportional to the overall change and the same for all individuals – the earnings elasticity of individuals to aggregate GDP changes is thus homogeneous across the population. However, if aggregate GDP changes affect different types of agents differently, and these differences are also systematically related e.g. to different consumption responses of these individuals, then the results of the representative agent model need to be adapted to this. The object describing earnings elasticities to GDP can be defined as:

$$\gamma_h = \frac{dE_h}{dY}\frac{Y}{E_h} \tag{1}$$

Where γ_h is the earnings elasticity of an individual household or some selected group of households, e.g. an income decile, E_h are the earnings of this group or individual household, and Y is aggregate GDP in the economy. Werning (2015) and Bilbiie (2008) explore theoretically how differences in individual earnings responses to aggregate changes affect the representative agent model. Werning (2015) discusses the implications of varying assumptions on the earnings elasticity of different households to aggregate output changes for equilibrium macroeconomic outcomes under incomplete markets. Incomplete markets matter because they prohibit individuals from insuring themselves to shocks. After the shocks, e.g. to their income, have realised, individuals who were unable to insure themselves become unequal to other individuals. This is how inequality materializes.

Werning (2015) shows the conditions under which, even with incomplete markets in the form of lacking liquidity and therefore borrowing, a model with idiosyncratic income uncertainty and borrowing constraints can lead to the same results as a representative agent model – it requires that individual income moves proportionally with aggregate income. The homogeneity of earnings elasticities to aggregate changes leads to an „as if" result – even though there are different households, with different levels of income, the representative agent results survive if the sensitivity of all agents to aggregate income changes is the same. Under this „as if" set-up, the level of consumption is higher with uncertainty due to precautionary motives, but the elasticity of aggregate consumption to interest rates is not affected – it is the same as in the RA case. Therefore the aggregate dynamics of the model are unchanged e.g. after a monetary policy change, although we have different agents with different income levels and incomplete markets cannot insure them against these income shocks. In the case of borrowing and positive liquidity and under logarithmic utility, the equilibrium outcome can again be aggregated when individual incomes are proportional to aggregate income. However Werning (2015) shows that when we depart from the assumption that individual income changes proportionally to aggregate income and if uncertainty is countercyclical then aggregate demand becomes more sensitive to interest rates - the heterogeneous agents model and its results are therefore different from the representative agent one. The result is shown by adopting a set-up with household unemployment risk and a varying extensive margin for employment. In this framework, aggregate output both affects positively the income of the employed, but also reduces the risk of being unemployment or underemployed and therefore also lowers uncertainty and precautionary savings motives. In essence, Werning (2015) provides the hypothesis tested in this paper – whether earnings elasticities to aggregate GDP of different groups are homogeneous or heterogenous. And furthermore, if they are heterogeneous what are the macroeconomic implications of that.

To answer the first question on the earnings sensitivity of different income groups, the literature has taken two approaches. One of the approaches takes granular administrative tax data to explore income risk – how are different incomes at the individual level reacting to economic growth. Guvenen et al. (2017) explore the most detailed dataset on earnings volatility in the United States, documenting more than three decades of individual income changes based on an administrative micro dataset. Using this large dataset, the authors estimate a regression on how the individual income of workers change in relation to aggregate GDP changes. The coefficient from this regression is then defined as “workers beta”.

Guvenen et al. (2017) document a number of important stylized facts about the relationship between earnings growth and GDP growth. First, the authors find that earnings risk, estimated as earnings elasticity to GDP fluctuations, is U-shaped in relation to the level of earnings. The earnings elasticity is very high at the beginning of the income distribution and decreases along the income distribution, with the exception of workers at the top of the distribution where it increases again. Lower-income households are more affected by recessions because their labor income is more sensitive to cyclical downward shocks. The earnings elasticity for the lowest income percentiles is above 3.5, while it decreases to a little above 1 across the middle of the distribution. This downward pattern can be explained by the fact that the drop in the income of some groups such as low-wage earners and the young is much bigger e.g. during recessions – these groups are much more sensitive to aggregate output fluctuations. On the very right end of the distribution, for the extremely high earners, it increases again. The high earners tend to be more exposed to the business cycle as a whole as they are affected mainly by profits and capital income. This is in line with some studies which concentrate on top earners - Parker and Vissing-Jorgensen (2009) e.g. report that income is mostly affected by the business cycle at the top of the income distribution. The differences in the sensitivity of earnings across households also means that different macroeconomic stabilisation policies with different targeted households tend to have important distributional, but also aggregate implications.

Due to their detailed dataset, the authors also can decompose the sample into separate groups and explore their income sensitivity. Men are more exposed to this aggregate risk and younger workers tend to be more exposed to the risk than old workers. Around the middle of the income distribution the groups with the highest earnings risk are male, young and workers in construction and durables manufacturing. Workers that work for large employer tend to have a lower sensitivity of their income to aggregate GDP fluctuations. Age also seems to significantly affect workers betas – while for the youngest age group 26 – 35 years, the estimated earnings elasticity is always above 1 and reaches up to 3 at the two ends of the income distribution, for the oldest category – 56 – 65 years, GDP beta is below 1 above the 30th percentile and below the 95th percentile and reaches very low levels of close to 0 in the middle of the distribution.

Detailed administrative data to track the individual income history is the most useful source to estimate income risks and income elasticities, yet it is very difficult to obtain. Another approach therefore has explored aggregated data from the income distribution to obtain the earnings elasticity of specific income groups or at specific points of the distribution. Dossche and Hartwig (2019) analyse the relationship between business cycle fluctuations and income changes of individuals and households in the four biggest economies of the Euro area. The study explores data from the EU Statistics on Income and Living Conditions (EU-SILC) and more concretely longitudinal individual income data to calculate the earnings elasticities in Germany, France, Spain and Italy for different parts of the income distribution. The authors again report worker betas following Guvenen et al. (2017) by calculating the earnings elasticities of workers in relation to GDP growth and find that these are higher for the lower income bins, especially so in Germany, France and Italy. For Germany and France, lower income households have an elasticity of around 1.4, but this decreases along the income distribution, while in

Italy the lowest quintile of income has an earnings elasticity of above 2.0, while the two top income quintiles - below 0.5. In Spain they report relatively similar, but high elasticity of around 1.4 across the whole income distribution. In some sense, household income risk according to these results behaves in the Euro area similarly to the findings about the United States from Guvenen et al. (2017) – a clear downward pattern can be identified. The results are however different in magnitude and this is expected – in the first case, administrative data provides us with the full information on the income mobility of individuals. In the second case, aggregated data averages out large income increases and decreases – as it measures the aggregated average incomes of a specific part of the population, e.g. the lowest decile, or at a specific point, e.g. at the 10th decile. The comparison between the two approaches therefore is useful to identify the pattern of how workers betas change along the income distribution. This is also the main goal of this paper.

Dossche et al. (2021) similarly report worker betas, defined as the elasticity of labour income relative to aggregate GDP growth changes, for the Euro area as a whole by using data from the biggest 4 economies in the Euro area. The results for the Euro area again show that lower-income households are more sensitive to changes in GDP growth – while for the lowest-quintile household earnings elasticity to GDP changes is around 1.5, for the higher quintiles it is slightly above 0.5. The Euro area as a whole, described by the Member states for which data is available, is also the focus of this paper so the results of Dossche et al. (2021) are directly comparable to some of my results.

To answer the question on the earnings elasticity of different agents in the most precise way therefore the researcher needs detailed microdata following the same agents wherever possible. Going further and merging different administrative datasets, Patterson (2022) quantitatively examines both the question of the earnings elasticity of different groups to aggregate shocks and the importance of this object for macroeconomic outcomes and macroeconomic policy. In a broad set of HANK models, aggregate consumption responses can be amplified if the same households or individuals which are mostly susceptible to an aggregate shock are also the ones with the highest MPCs. Patterson (2022) first measures the MPCs of different socioeconomic groups and then quantifies the unequal incidence of aggregate shocks on the individual income of those groups. To do so, the paper merges two datasets – the US Panel Study of Income Dynamics (PSID) and the Longitudinal Employer Household Dynamics data (LEHD) to obtain a detailed panel dataset on household earnings history, income and consumption. Using this dataset, the MPCs of different groups, grouped by age, race and income, can be estimated and are documented to be large for some groups and very heterogeneous, as discussed in previous studies (Japelli and Pistaferri, 2010). More importantly, the groups with the highest MPCs according to this estimation, are the groups which have the highest elasticity of their income to aggregate GDP fluctuations. This important covariance between earnings elasticity and MPCs is then also shown to hold with data from other sources – e. g. the Italian Survey of Household Income and Wealth (SHIW). Second, using geographical variation and data on commuting zones in the US and their relative economic performance during recessions, Patterson (2022) shows that the areas with higher covariance between MPCs and earnings elasticity experience more volatile business cycles – they have stronger economic boom periods and deeper economic downturn periods. This serves as evidence that the amplification mechanism proposed by Patterson (2022) indeed holds in the aggregate.

This paper serves as a contribution to each of these approaches. First, I evaluate the earnings elasticity for different income deciles in the Euro area using two sources of aggregated income data – the

Distributional National Accounts Data from the World Inequality Database (WID) and the Global Repository of Income Dynamics (GRID). These estimates can therefore be directly compared to the ones by Dossche et al. (2021). Second and more importantly, I construct a new dataset of earnings changes using the panel part of the ECB Household Finance and Consumption Survey (HFCS). Part of the households surveyed in this ECB household survey are repeated across waves and I use this to construct the after-tax earnings changes of these households. By using this new dataset, I am able to evaluate at the microlevel the earnings elasticity to aggregate GDP changes for the countries in the Euro area similarly to the way that Guvenen et al. (2017) does it for the United States. Finally, using my newly constructed dataset of earnings changes, the estimated earnings elasticities and a unique question from the HFCS on self-reported MPCs of households, I can explore the relationship between households' earnings elasticities and the MPCs of households. Furthermore, in line with Patterson (2022) I evaluate quantitively how important this mechanism is as an amplifier of macroeconomic dynamics.

III. Methodology and Data

The sensitivity of individual earnings to aggregate shocks is a key moment for any model which takes household heterogeneity into account. Empirical evidence on this question is however scarce because the necessary data has been hard to obtain. To estimate these earnings elasticities of different household groups the researcher needs a panel following the same households through periods of changing economic conditions. Using decades of US administrative tax data, Guvenen et al. (2017) report estimates on the earnings elasticities of households to aggregate output changes. Their main baseline regression is given in Equation 1:

$$\Delta log E_{h,t} = \beta_0 + \beta_1 \Delta log Y_t + \epsilon_{h,t} \tag{2}$$

The estimation thus delivers the elasticity of the earnings of individuals workers groups by estimating the coefficient β_1 – so called "workers betas" or in our case households betas. I estimate and report these coefficients in Section 4 for a number of different specifications.

The most straightforward and precise way to estimate this question is by using a similar household panel with data on earnings, where the evolution of individual earnings can be tracked throughout a number of decades and the response of individual earnings to aggregate GDP changes can be evaluated. Such data is not available for the Euro area. To estimate the patterns of earnings elasticities therefore, I take a two-fold approach. First, I use aggregated data on earnings by different income groups, I analyse the earnings changes of different deciles and estimate how they react to aggregate income changes. While this data does enable us to track the dynamics of income of the same households throughout time, it nevertheless enables us to estimate the earnings elasticities of the given income groups. In terms of aggregate outcomes this information is still relevant, since the sum of the different income groups aggregated earnings equal total income in the economy. Secondly, and most importantly, I also build a new dataset of earnings dynamics. I construct this dataset by using the panel component of the ECB HFCS dataset. The ECB HFCS contain a repeating panel component for around half of the respondents in some of the countries of the sample. By merging the individual waves and correcting manually for taxes, I build a dataset of after-tax income changes for a representative sample of households in around 10 countries, including the big four countries – Germany, France, Italy and Spain. The composition of countries changes from wave to wave as some countries have introduced their panel component of the survey only recently.

Using this newly created dataset, I can also answer the further question on the relation between the earnings elasticity and the MPCs of low and high income groups. Patterson (2023) is the first to study this question in relation to both the earnings elasticity and its covariance with household MPCs by using administrative US data. Patterson (2023) finds that the mechanism – that workers with high MPCs often have jobs which are highly impacted by downturns - is quite strong in the US and contributes to the amplification of recessions. In the baseline specification, Patterson (2023) estimates:

$$\Delta log E_{h,t} = \beta_0 + \beta_1 MPC_h + \beta_2 MPC_h \times \Delta log Y_t + \epsilon_{h,t} \tag{3}$$

Where Δ log $E_{h,t}$ is the change of income of household h throughout two consecutive measurements, MPC is the self-reported household response to the MPC question in the survey, and Δ log Y_t is the 3 - change in GDP throughout the same period. The coefficient β_2 is the coefficient of interest for us.

It enables us to then evaluate how the amplification of shocks through the MPC is affected by the unequal incidence of aggregate shocks by quantifying the covariance in this expression of the aggregate MPC in the economy:

$$MPC = \sum \frac{dC_h}{dE_h}\frac{dE_h}{dY} = \sum \frac{dE_h}{Y}\frac{dC_h}{dE_h} + \left(Cov\frac{dC_h}{dE_h}, \gamma_h \right) \tag{4}$$

where C_h is the consumption of household h, Y is aggregate output, and γ_h is the elasticity of household earnings to aggregate output. While the first term gives us the aggregate MPC, the second term decomposes it in the weighted average MPC and the covariance between the MPC and earnings elasticity. As noted in Section 2, this object γ_h, which we aim to estimate in this project, is given by $\gamma_h = \frac{dE_i}{dY}\frac{Y}{E_i}$. Bilbiie (2008) and Werning (2015) explore the theoretical implications of different values for γ_h, without empirically evaluating it.

As discussed, I approach the issue of how to measure the earnings elasticity of different income groups and their systematic differences in the MPCs through three separate data sources and compare the results. First, I use data on income changes for income deciles from the Distributional National Accounts (DINA) developed as part of the World Inequality Database (WID). This data does not enable me to estimate the earnings elasticity for a panel of individual households, as one cannot follow the same households. I can rather follow and evaluate the earnings elasticity for aggregated parts of the distribution. With aggregated data it is feasible to use only the percentiles or deciles of the income distribution and estimate the earnings elasticity of each percentile or decile to overall aggregate GDP fluctuations. The DINA data on household earnings are grouped by decile and run from 1980 to 2020. Secondly, I use data from the Global Repository of Income Dynamics (GRID) database. This database explores income risk in advanced economies by gathering administrative tax data in a comparable way. In this dataset income changes for different parts of the income distribution are reported as aggregated average values.

Finally and most central to this paper, I use the ECB Household Finance and Consumption Survey. The ECB Household Finance and Consumption Survey (HFCS) is a detailed microdata survey, conducted every 3 years, where households in all Euro area countries respond to an extensive questionnaire regarding their household characteristics, income, consumption, assets and liabilities, as well as further information. I construct a new dataset by merging the four available waves of the HFCS and computing earnings changes for any households which is included in two consecutive survey waves. The HFCS has a panel dimension in a changing number of countries throughout different surveys. For a full list of the countries and the years of the panel component, see Annex I. I therefore build a dataset of changes in household total earnings of households between two waves – between 2010 and 2014, between 2014 and 2017 or between 2017 and 2021. For my baseline regression, to calculate the households' betas, I therefore have compiled the data on earnings changes, combined with data on aggregate output changes.

By using three different datasets I deliver two important results. On one hand, I estimate the same regressions across all three different data sources and compare the patterns of the results. On the other hand, the two specific types of data estimate two different types of earnings elasticities. The DINA and GRID data, unlike the HFCS data, enable us to make only estimations based on the aggregated averages per decile or percentile. These aggregated data hide the two channels of income changes and employment changes, but also do not take into account mobility along the earnings distribution. Generally, the question on the relative earnings changes given aggregate output fluctuations has an intensive and an extensive margin. Business cycle upturns and downturns have two elements that determine workers income – they change their real earnings, but also enhance or worsen their employment chances and risks. As an approximation, the difference in magnitudes of workers betas between the two types of data can be perceived as coming from the different level of mobility workers experience and as coming from extensive and intensive margin dynamics. Using the HFCS I estimate earnings elasticities also by aggregating the data, similarly to the other two datasets. This enables me to show the additional information gained by using individual level, panel data following the same households throughout time.

IV. Results

I now pursue to estimate Equation 1 and therefore to obtain the coefficient β_2 – the coefficient of workers betas with the three different datasets.

1. Distributional National Accounts Data

First, I use the Distributional National Accounts (DINA) data available from the World Inequality Database (WID). In this estimation, I follow the annual changes in the aggregated earnings of the ten income deciles and evaluate how they vary with aggregate output fluctuations for the period since 1995, although longer horizons are also available for some countries. I estimate the earnings elasticities for the biggest four Euro area economies – Germany, France, Italy and Spain, which together cover more than 70% of Euro area GDP, which I consider a good enough approximation to define this measure as my Euro area measure. I estimate equation 5, where $E_{d,t}$ are aggregated earnings for each of the ten deciles and Y_t is real GDP:

$$\Delta log E_{d,t} = \beta_0 + \beta_1 \Delta log Y_t + \epsilon_{d,t} \tag{5}$$

Figure 1 presents the results from the estimations using the DINA data. I estimate Equation 4 using pre-tax data on incomes of households in each of the ten deciles for the data sample consisting of the big four countries. The lowest decile has the highest sensitivity to aggregate output of above 1.5. The deciles in the middle of the distribution have an earnings elasticity of 1 and therefore seem to move one-to-one with the changes in aggregate economic output. At the very top, the top decile has again a slightly higher value for beta of around 1.3. The pattern of workers betas obtained therefore is similar to Guvenen et al. (2017). For most of the income distribution a clear downward pattern emerges – the lower income groups have higher earnings elasticities. Only at the top this pattern is reversed, which in the aggregate resembles a U-shaped pattern.

As a robustness check, I run the same regression but using aggregate income instead of aggregate GDP as economic output change variable. The results are very similar. In the above I use pre-tax income data since the after-tax income variable is not available for Italy and this biases my results and makes it difficult to compare with the results from the next section.

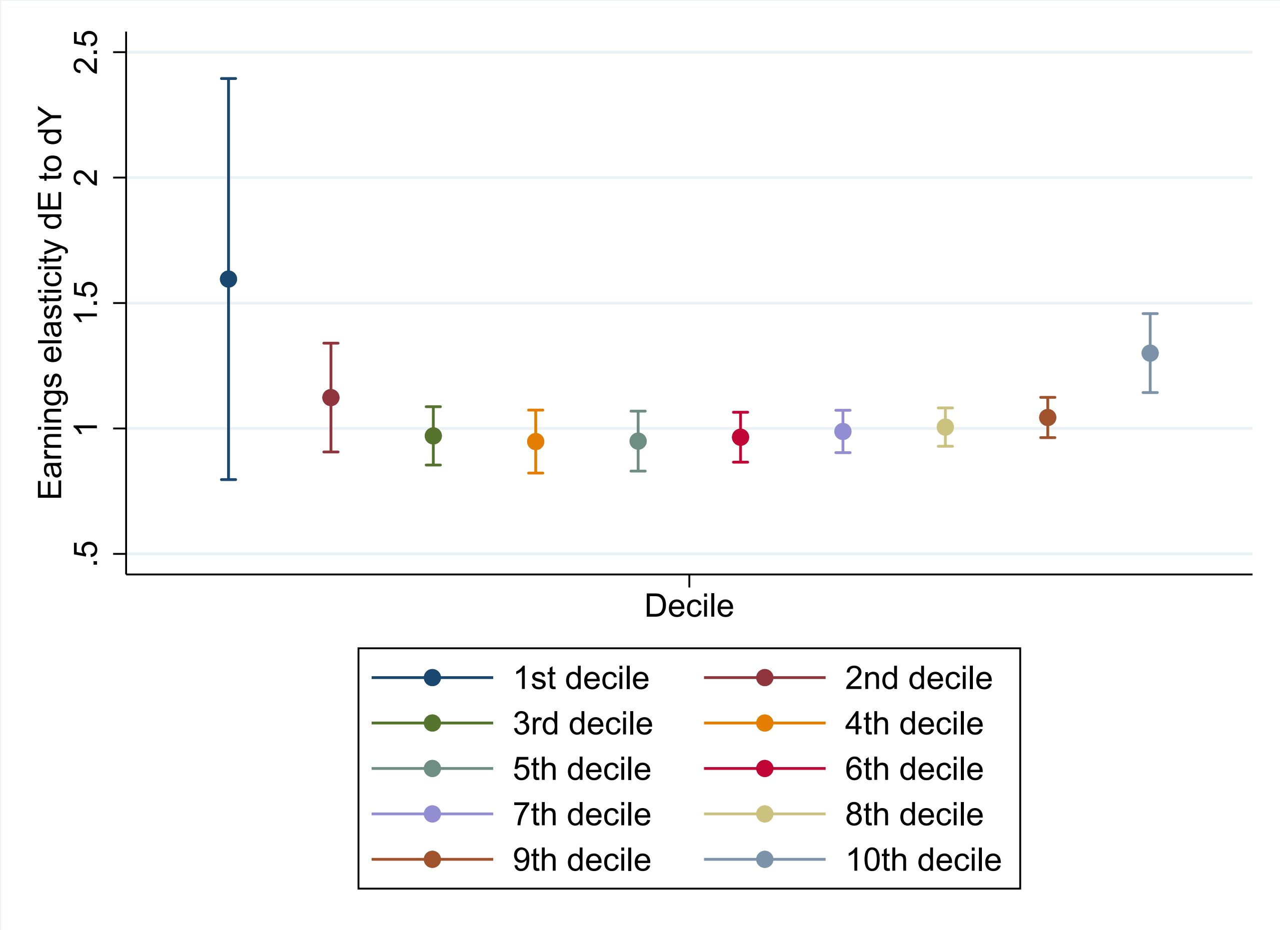

Figure 2: Earnings elasticity dE/dY using the DINA dataset for the Euro area
Notes: The estimation uses the sample of 1995 – 2020. Household income dE is pre-tax income, dY is the change in aggregate GDP.

2. GRID Data

Next I estimate a second set of earnings elasticities by income groups for the Euro area using the aggregated data from the Global Repository of Income Dynamics (GRID) database. The GRID database provides a set of micro statistics on income inequality and income dynamics at the individual level for a sample of countries in a harmonized way for cross-country comparisons. The database is gathered from national administrative records on earnings histories from each country. This large micro dataset is granular enough to provide precise measurements of specific moments also for subpopulations in each country – along age, educational attainment, income level and income shares. The panel character of the data then enables me to track the history of income dynamics for those specifics groups or points of the distribution throughout time. For the purposes of my study, this provides me with the time series of how the incomes of different income groups evolve over time and again enables the analysis in terms of earnings elasticities.

Out of the available 13 countries, in the GRID database, 4 are from the Euro area, and are again the big four countries – Germany, France, Italy and Spain. I use the data on the income level of selected income percentiles of the population (5th, 12.5th, 25th, 37.5th, 50th, 62.5th, 75th, 87.5th, 95th and 99th percentiles) for all available years in each country [2]. Using this data on the percentiles of the

[2] The GRID database does not report income per decile as in the WID database, but rather selected percentiles across the income distribution, so I choose 10 percentiles starting from the very low (5th) and going up to the highest (99th).

log earnings distribution, I compute the annual earnings changes for each of these percentiles. I then regress again these earnings changes to real GDP changes (Equation 6), where $E_{p,t}$ are aggregated earnings for each percentile and Y_t is real GDP:

$$\Delta log E_{p,t} = \beta_0 + \beta_1 \Delta log Y_t + \epsilon_{p,t} \tag{6}$$

I estimate Equation 6 using the reported earnings of households at each of the selected percentiles. The sample covers different periods in each country between 1995 and 2016 as data availability varies.

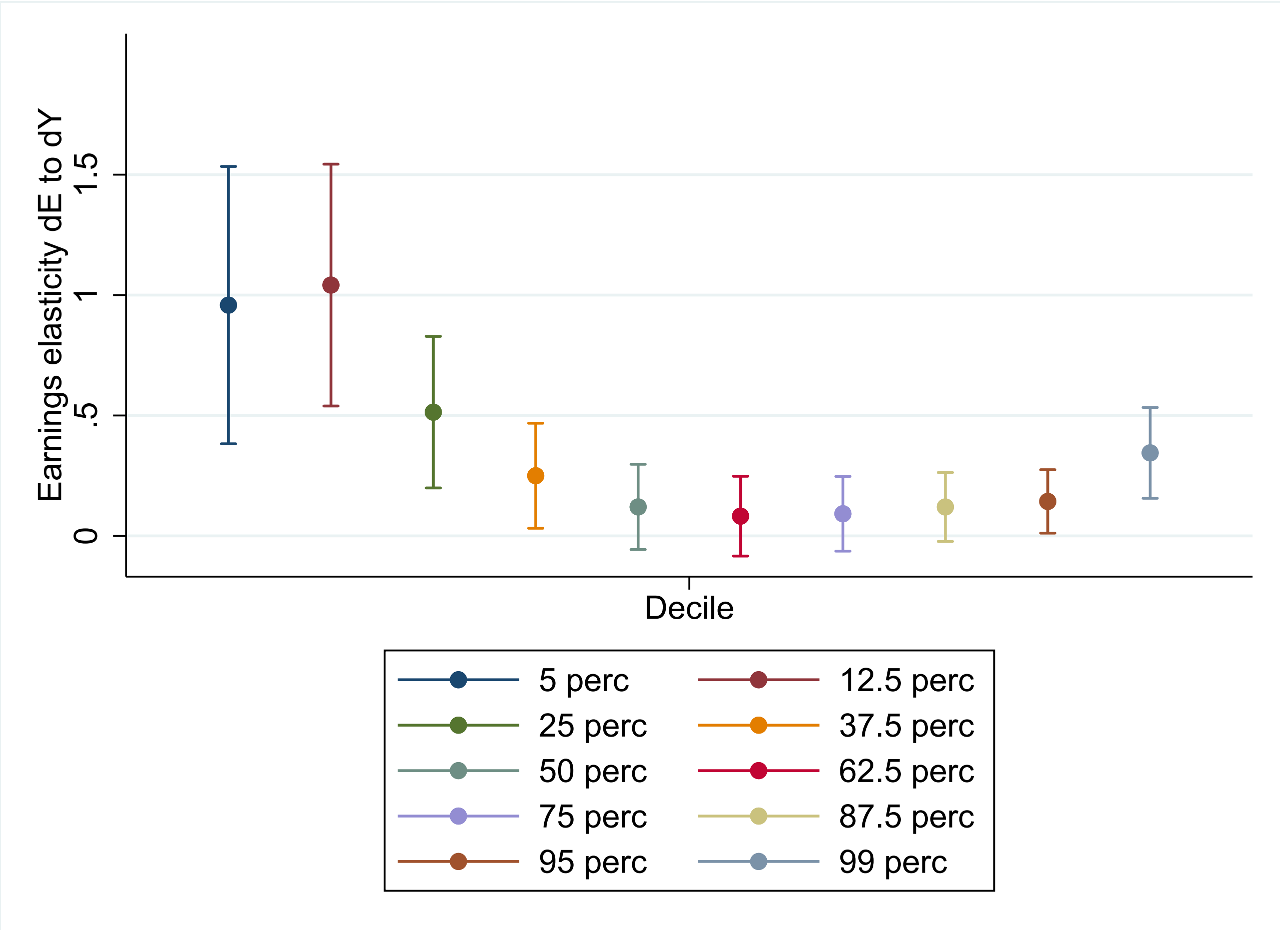

Figure 3: Earnings elasticity dE/dY using the GRID dataset for the Euro area
Notes: The estimation uses the sample of 1995 – 2016, but the country availability of data varies. Household income dE is pre-tax income at a specific point in the distribution - the given percentile, dY is the change in aggregate output.

Figure 3 presents the results from the estimations using the GRID data. The pattern is very similar to the one estimated using the DINA data. It is also similar to the results by Guvenen et al. (2017). Earnings elasticities have a downward sloping trend up until the highest percentiles. The lowest deciles have the highest sensitivity to aggregate output of around 1. Around the middle of the distribution the magnitude is much smaller and the earnings elasticity is less than 0.20 The U-shaped pattern can then be identified only at the very right end of the distribution and is only present because we take the highest earners into account as well – the earnings elasticity increases from 0.14 to 0.34 between the 95th and the 99th percentile. The magnitudes here are only partly comparable with the previous

These points of the distribution are assumed to serve as an approximation for each decile estimation - since my GRID data measures income changes at a point of the income distribution and not for the whole bin, they are especially susceptible to the problem of mean reversion. However they are also informative as each of these data points can be taken as the median for a specific bin or decile, which then moves proportionally with the change of this data point. The inability to track a specific household, but rather a point of the distribution, which is represented in each wave by a different household, provides us therefore very smoothed information on the dynamics of the whole income distribution and its bins. This smoothing of the elasticities provides us with estimates which are much lower in magnitude. Even so, the pattern of decreasing earnings elasticites across the income distribution persists.

3. HFCS Data

Finally, I turn my attention to a source of detailed microdata on household income, consumption and financial assets – the ECB Household Finance and Consumption Survey. The survey provides me with granular data on the balance sheets and earnings of households in all Euro area member states for each of the four waves of the survey, which have been conducted approximately every three years. The HFCS includes a panel component in some of the countries, where some households interviewed in the previous wave are also interviewed in the next one.

I make use of this panel feature of the HFCS by constructing a new dataset to follow income dynamics for households in the countries with a panel component. I pursue in the following way. First, I merge all four waves and identify those households which are included in two consecutive datasets. I link them by using their past household and individual ID. Secondly, I clean the dataset for any irregular income reportings (e.g. annual income of 1 EUR or similar assumed mismeasurements due to respondents misperceiving the question). Thirdly, similar to Guvenen et al. (2017) I will be interested in the after-tax earnings of households, since these are the ones that have an implication for their disposable income, their consumption patterns and dynamics and therefore for any macroeconomic dynamics. The HFCS does not report after-tax income, so I construct it by using income tax brackets reported by the OECD for each country. I therefore approximate and impute after-tax earnings and income for each household. This computation does not include the full variety of country-specifics tax rules and possible tax exemptions, but should serve well enough as an approximation for after-tax income.

Figure 4 reports the household real income earnings changes from my newly constructed database. To compute the real income changes, I deflate the second observation on income for each panel household with the HICP inflation gathered between the two waves. Even though the HFCS takes place normally every three years (with the exception of the pandemic induced delay in the fourth wave), the reference period is specific for each country as the period in which the interviews are conducted varies considerably between countries. In addition to that, the reference period for the question on income varies between the last 12 months and a specific year. I take this into account and compute the inflation accumulated between the end of a given year and the end of the next reference year for each country reported in Table A1 in the Annex. Similarly, I compute GDP changes as the accumulated GDP growth between the end of the first and second year used as reference period between the two corresponding waves. These GDP changes for my dataset are reported in Figure 5.

Having constructed this new dataset I can calculate the income changes for all households in the dataset. I obtain the annual changes in the earnings of households, which are present in two consecutive

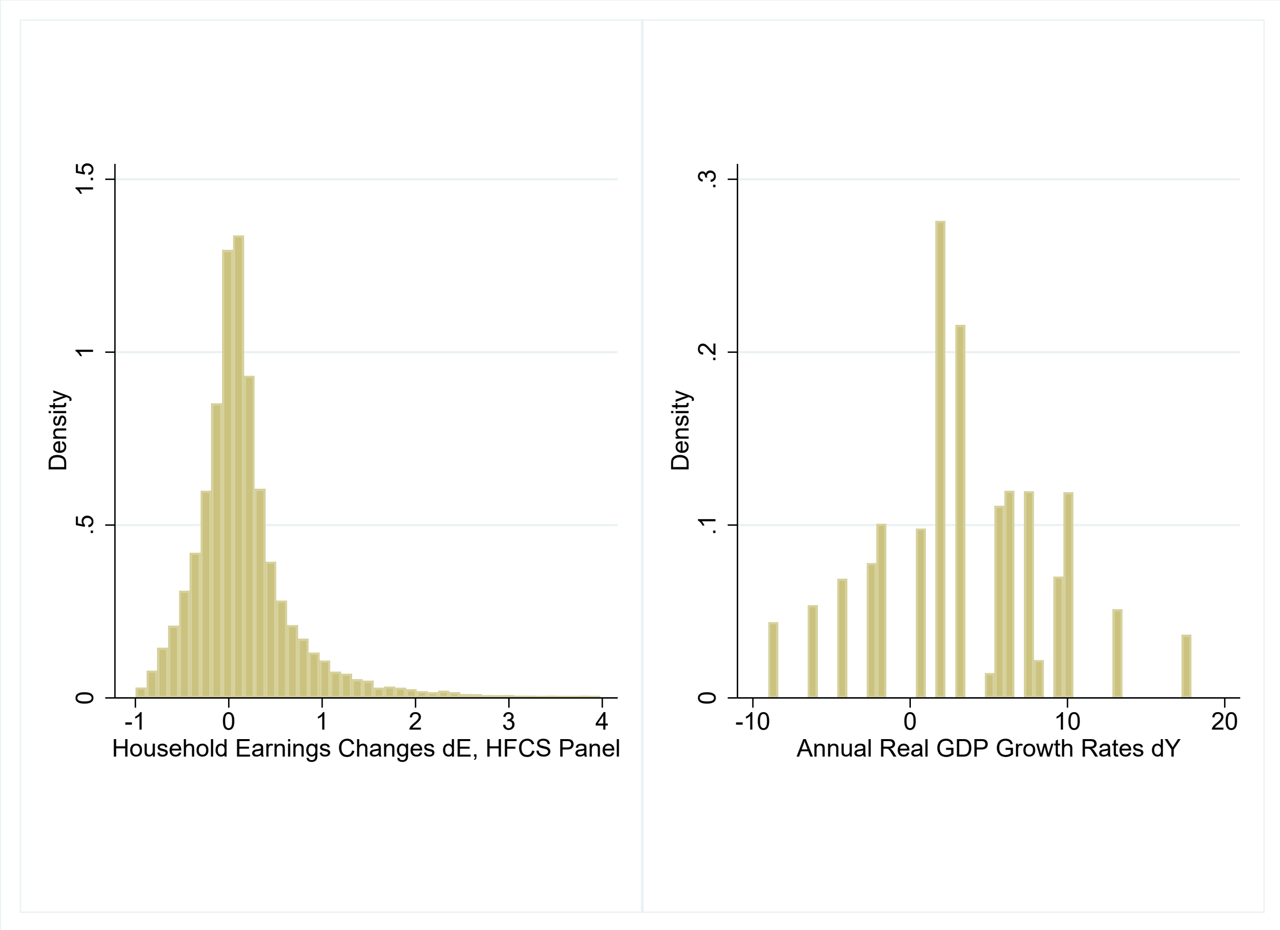

Figure 4: Descriptive statistics on earnings and aggregate GDP changes.
Earnings changes reported exclude extreme outcomes - positive and negative income changes of more than 400%.

waves of the HFCS. I group households in ten income deciles in each country, based on their starting level of income. I then estimate again the following equation for each decile separately, where E_h are individual household earnings and Y_t is real GDP:

$$\Delta log E_h = \beta_0 + \beta_1 \Delta log Y_t + \epsilon_{p,t} \tag{7}$$

I therefore obtain earnings elasticities given by β_1, the workers beta, for each income decile. In this estimation, I define these earnings elasticities to be for the Euro area, as I cover between 9 and 11 countries depending on the wave, which constitute more than 70% of the Euro area GDP.

Figure 6 presents the results from the estimations using the HFCS data and the available Euro area countries in the HFCS and using total household earnings as the earnings variable at the household level. The lowest decile has a high sensitivity to aggregate output of 3.40, which is very close to the estimate for the same group by Guvenen et. al (2017). Earnings elasticities very smoothly decline across the income distribution, with the next income deciles having an elasticity of between 3 and 2. The deciles in the middle of the distribution have lower earnings elasticity of around between 1.4 and 1.2. For the two highest income deciles the earnings elasticity decreases further and is below 1 - at 0.8 and 0.6 for the ninth and tenth percentile respectively.

Generally, I recover the similar pattern identified in the previous estimations – earnings elasticities decrease along the income distribution, yet this time they do not follow a U-shaped pattern and do not increase at the top end of the income distribution. This is partly explained by two factors – first there is top coding and well-known underreporting of the top incomes in income surveys. This means that my income change measurement for the top decile is probably biased downwards. Secondly, both from the results from the estimation with the GRID data and from Guvenen et al. (2017), is is clear that the U-shaped pattern is driven by the extremely high incomes. In my estimations above using GRID data, only the 99th percentile of the income distribution shows again an increased earnings elasticity in comparison to the middle-income groups. In this estimation using the HFCS data with ten deciles, the last decile is too wide. Extremely high incomes are only a part of it and do not drive enough volatility in income changes to increase the overall workers beta of this bin considerably.

My results are well in line with Guvenen et al. (2017) and are the right measurement of earnings elasticities, since they are derived from microdata. Both the pattern of decreasing earnings elasticities along the income distribution is in line with Guvenen et al. (2017) and with the above estimations using aggregate distributional data, with the exception of extremely high income groups. But more importantly, the magnitude of workers betas estimated from the HFCS microdata is also in line with the one reported by Guvenen et al. (2017) for US data.

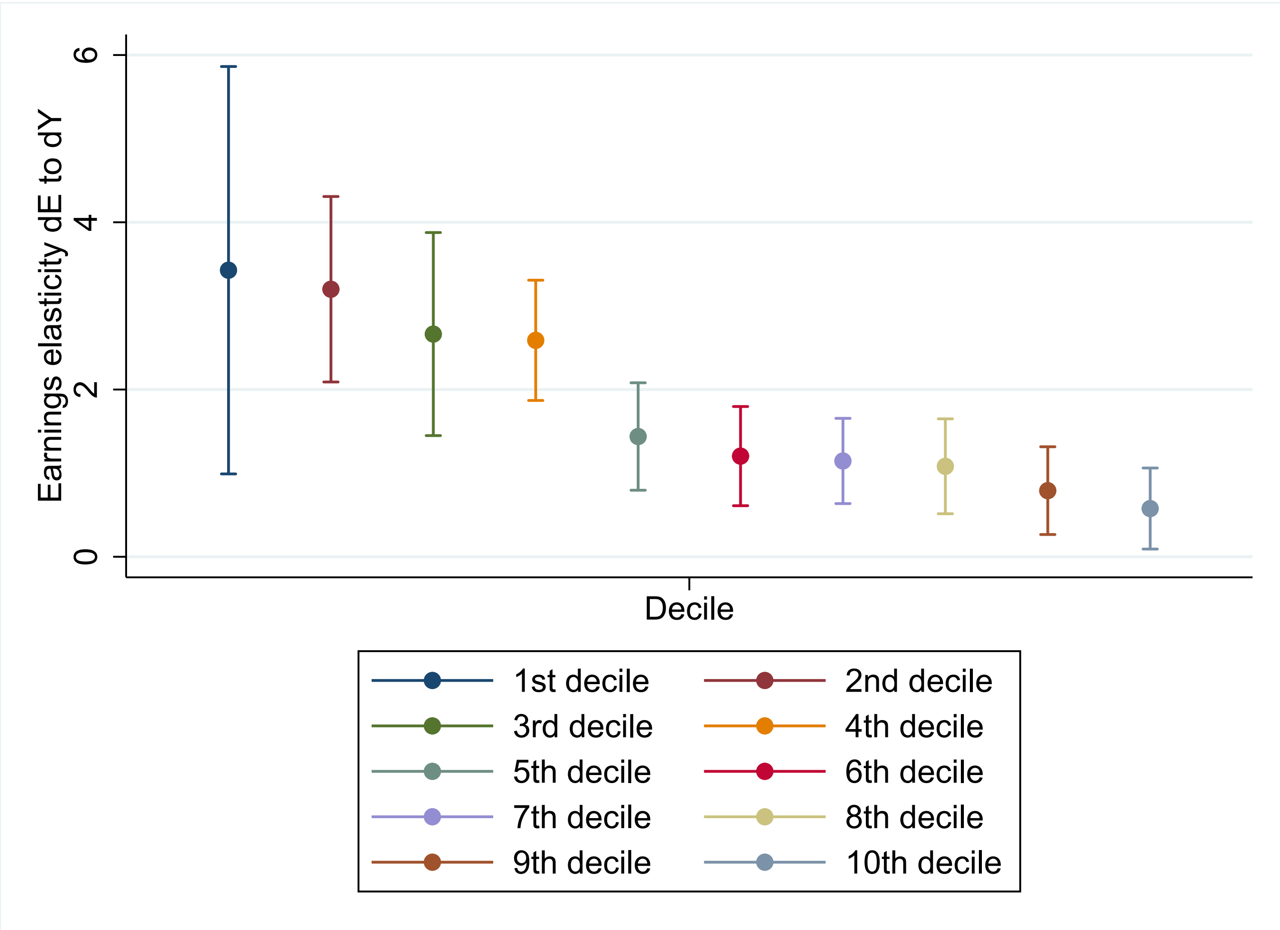

Figure 5: Earnings elasticity dE/dY using the HFCS dataset for the Euro area
Note: Earnings elasticities by decile using after-tax household earnings and excluding extreme observations where the rate of change was above 1000 percent (more than 10-fold increase or decrease in income).

In estimating the above results, I make one important adjustment. The HFCS reports standard household weights, which adjust the data so that the sample is representative for the population of the given country as is standard when working with houshold surveys. My estimation in (7) however aims at using cross-country variation of GDP growth rates between countries to estimate a hypothetical coefficient - the elasticity of each decile to aggregate output changes. If I use the initial household weights therefore my coefficient will be representative for the Euro area population, but the coefficients will be based in favour of larger countries, as their population will be overpresented and will have an overproportional effect in the estimation. From a theoretical point of view, it will be better if the estimated elasticity is not dependent predominantly on the larger countries. Therefore I construct new weights, which replicate the relative weight of each household in their population, but are scaled up or down so that all countries have the same population size and therefore influence on the final coefficient estimate. This enables me to deliver results which are not influenced predominantly by the larger countries in the HFCS.

4. Robustness

The results reported above are cleaned for extreme outcomes. Survey data from the HFCS is subject to considerable measurement errors, as data on income reported by individuals often is. The data includes a number of very extreme growth rates of incomes for a small number of individuals. They could bias the results. In the baseline estimation above therefore, I have dropped all observations, which register a change in their household earnings of more than 1000% (10 fold increase or decrease of income), as I consider these unrealistic and due to sampling error. Such restrictions are similar to those used in other papers in the literature, such as in Hendren (2017), where individuals with more than a threefold change in consumption of food are excluded, or in Gruber (1997), where observations with a greater than 1.1 log change in consumption of food are excluded. As a robustness check, I include also the extreme outcomes in my dataset and calculate again the same estimation for earnings elasticities. Figure 7 reports households betas for this full dataset. The results are very robust – the general decreasing pattern is similar to the results above, but the magnitudes are higher across the whole distribution. The earnings elasticity of the highest decile is now larger at around 4.9. It decreases smoothly across the income distribution to reach values of between 1.6 and 1.2 in the middle of the distribution and reaches a low of 0.80 for the highest income decile.

I conduct a further robustness check to evaluate the effects from aggregate GDP on household earnings by excluding from the estimation households who have been unemployed in one of the two periods where they are surveyed. Income shocks and income risks have in general an intensive margin, when households get an increase or a decrease in their income, and an extensive margin, if they become unemployed in one of the two observed periods. To test the importance of these two margins, I evaluate the above regression only for households which have been in uninterrupted employment throughout the sampling period. Figure 7 reports the results - they are very much in line with the estimates above. The intensive margin therefore drives the earnings elasticities estimated in our previous regressions, also due to the fact that households that were in constant employment constitute a very large part of our overall sample.

The three estimations for the Euro area show a consistent picture. The earnings of bottom income groups are more volatile to macroeconomic output fluctuations, while those in the middle vary less with the business cycle – the worker betas are the highest at the left end of the income distribution.

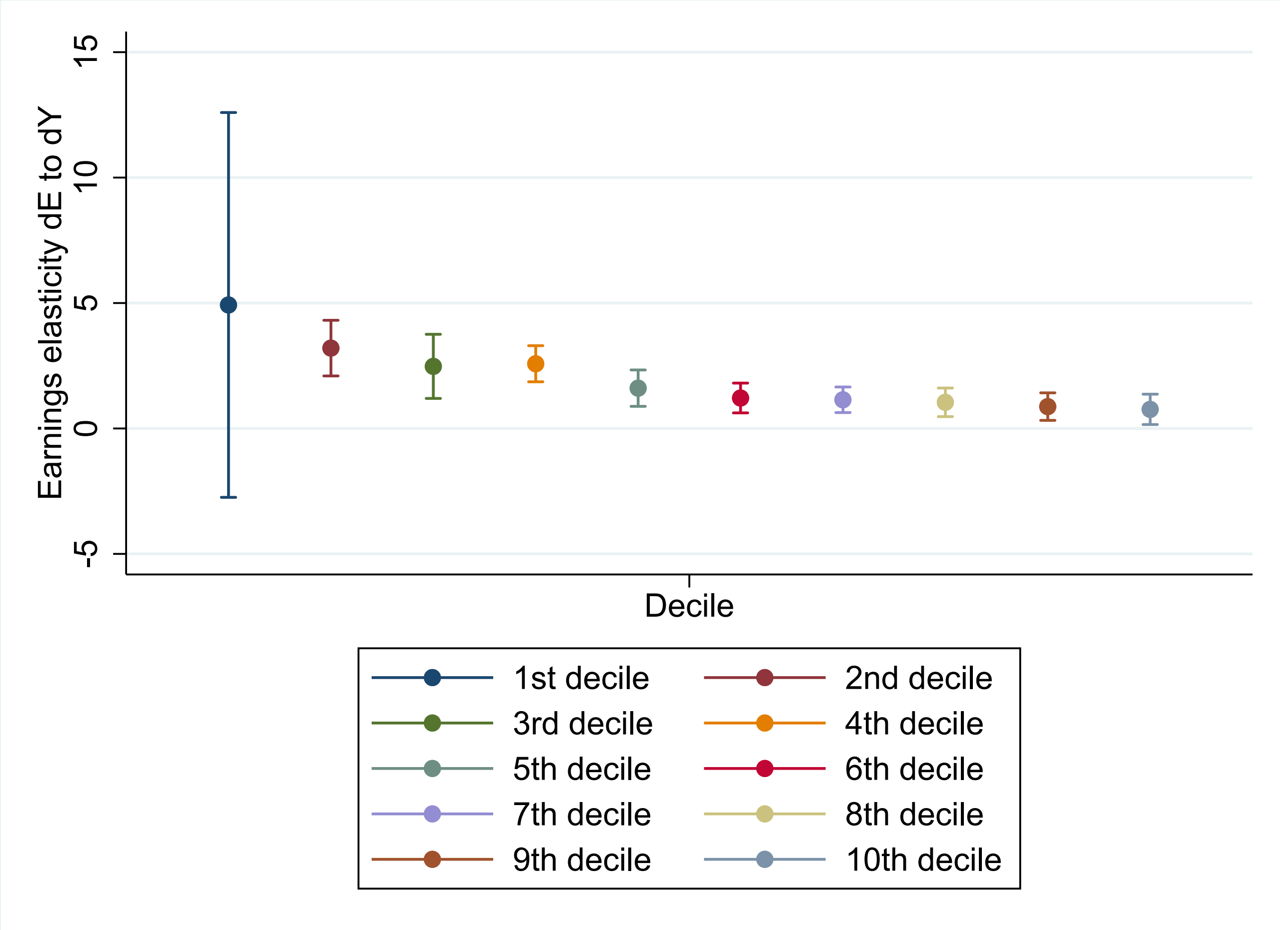

Figure 6: Earnings elasticity dE/dY by income decile, household income, Euro area, HFCS dataset
Note: The earnings measurement used here is total household earnings, full HFCS sample.

These results are in line with the detailed microdata evidence from Guvenen et al. (2017). The aggregated and microdata show a consistent pattern with the exception of very high incomes, but different magnitudes.

My results are the first to answer the question on the magnitude of earnings elasticity by using both aggregated data and microdata. The possibilities to compare results are therefore scarce. Dossche et al. (2021) however report workers betas for the Euro area using aggregated data. The study mainly explored how earnings inequality varies over the life cycle and reports the earnings elasticities for five quintiles in the Euro area using total disposable household income data (HY020) from EU-SILC. Income is grouped in quintiles and earnings changes are calculated in relation to the income in the initial period, whereas the initial period is the last two years. The earnings elasticity is then calculated as the elasticity of average income per quintile to aggregate GDP growth. Both the pattern is similar to my results – the lowest quintile is the most sensitive to aggregate fluctuations, but furthermore the magnitudes are similar. For the lowest quintile the authors obtain households betas of around 1.4, in line with my estimations from the WID and GRID data.

The estimations from my constructed dataset of earnings changes using the HFCS are also in line with the previous estimations using microdata by Guvenen et al. (2017). I argue that the panel data is the right measurement of earnings elasticities - aggregated data smooths the results as it does not take into account switching between deciles, but rather only calculates workets betas per the aggregated

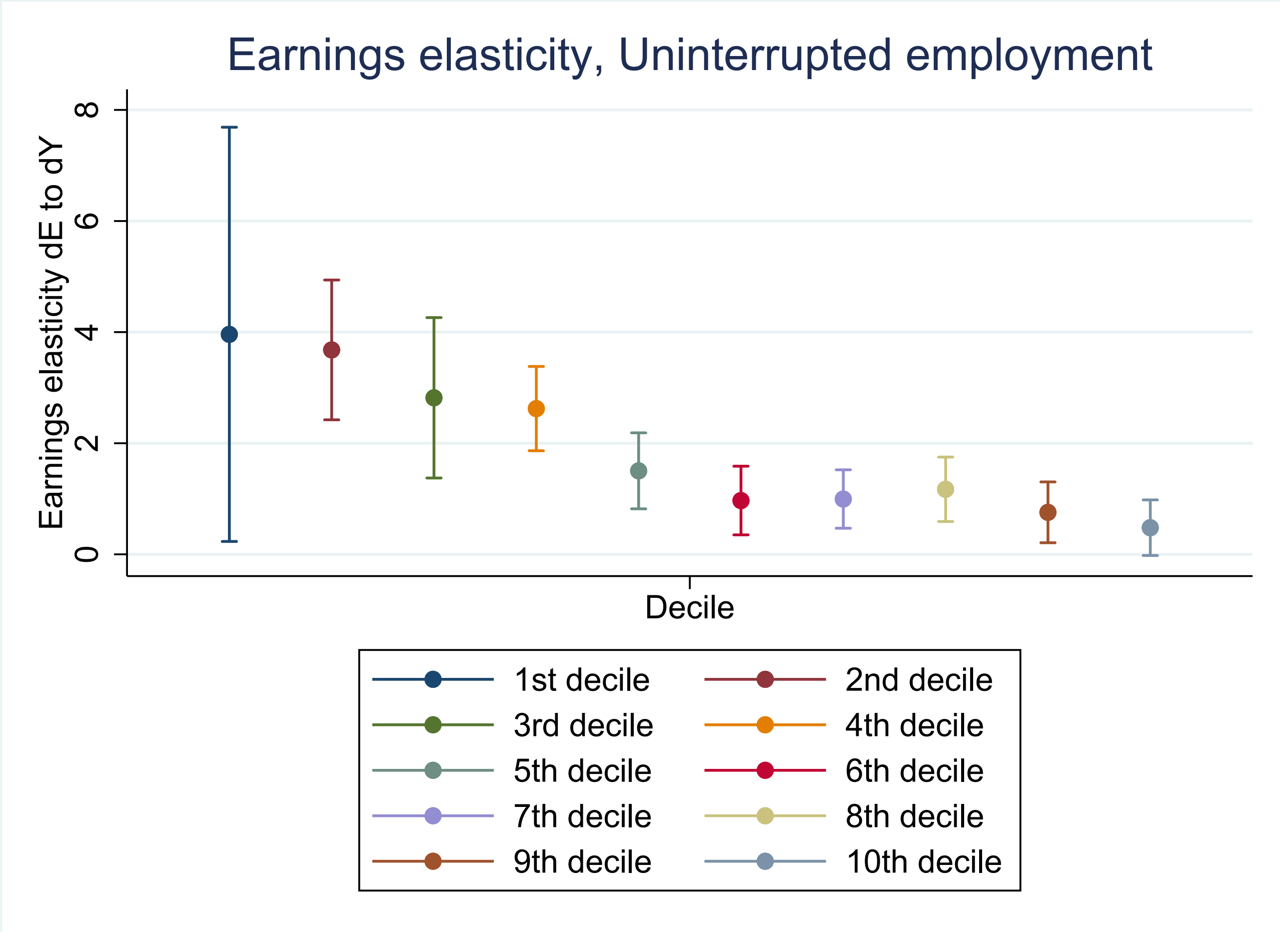

Figure 7: Earnings elasticity dE/dY by income decile for households without unemployment spells, household income, Euro area, HFCS dataset
Note: The earnings measurement used here is total household earnings, full HFCS sample for households which have been in employment in the waves where they are surveyed.

deciles or the income at a specific percentile. The HFCS results show again a downward pattern of earnings elasticities to aggregate GDP changes, but also the magnitudes are very much in line with the results by Guvenen et al. (2017) – the lowest income groups have households betas of above 3.4, whereas Guvenen et al. (2017) reports workers betas of around 3.7 at the left end of the income distribution.

5. Individual and Aggregated Level Data

The individual panel data from the HFCS enables me to follow earnings changes of households between different waves and therefore to account for the full mobility of households along the income distribution. The aggregated data from GRID and WID cannot produce the same results because the results there are automatically smoothed to the mean as the regressions are done at the group level - for specific deciles at a point in time, without explicitly following the same individual households. The magnitude for these aggregated regressions is automatically different as it excludes extreme changes which can happen at the individual level for idiosyncratic reasons. The aggregated groups ignore this mobility - an individual which started at the lowest decile in the HFCS can move up to the top decile and the HFCS estimations take this extreme change in earnings into account. The earnings changes of the aggregated groups do not reflect these changes - they only account for the average or median of the members of this specific income group at a point in time. How can we decompose and quantify the difference between the two measures of earnings elasticity and quantify the effects of this within group elasticity? The HFCS data enables me to do this.

Using the HFCS panel data, I construct a similar dataset to the one from GRID and WID and compare the estimates. I pursue in the following way. I group households along 10 income bins. I calculate the average earnings of each of these bins for each wave and the earnings changes between waves. This enables me to estimate again Equation 5 with earnings changes by income groups from the HFCS. I then compare the results from these aggregated earnings elasticities from the HFCS with the individual earnings elasticities obtained in Section 4. For the HFCS individual earnings elasticities, I compute a quadratic fitted line of the ten points estimated above in Section 3. Figure 8 plots these fitted predictions of the individual and aggregated HFCS earnings elasticities, as well as the share of the aggregated earnings elasticities to the individual level earnings elasticities. This enables me to analyse how much of the overall earnings elasticities estimated at the individual level can be explained by upward or downward mobility of individuals. As is obvious, this is a phenomenon that is most important at the left side of the income distribution. It then becomes less and less important at the high end of the distribution.

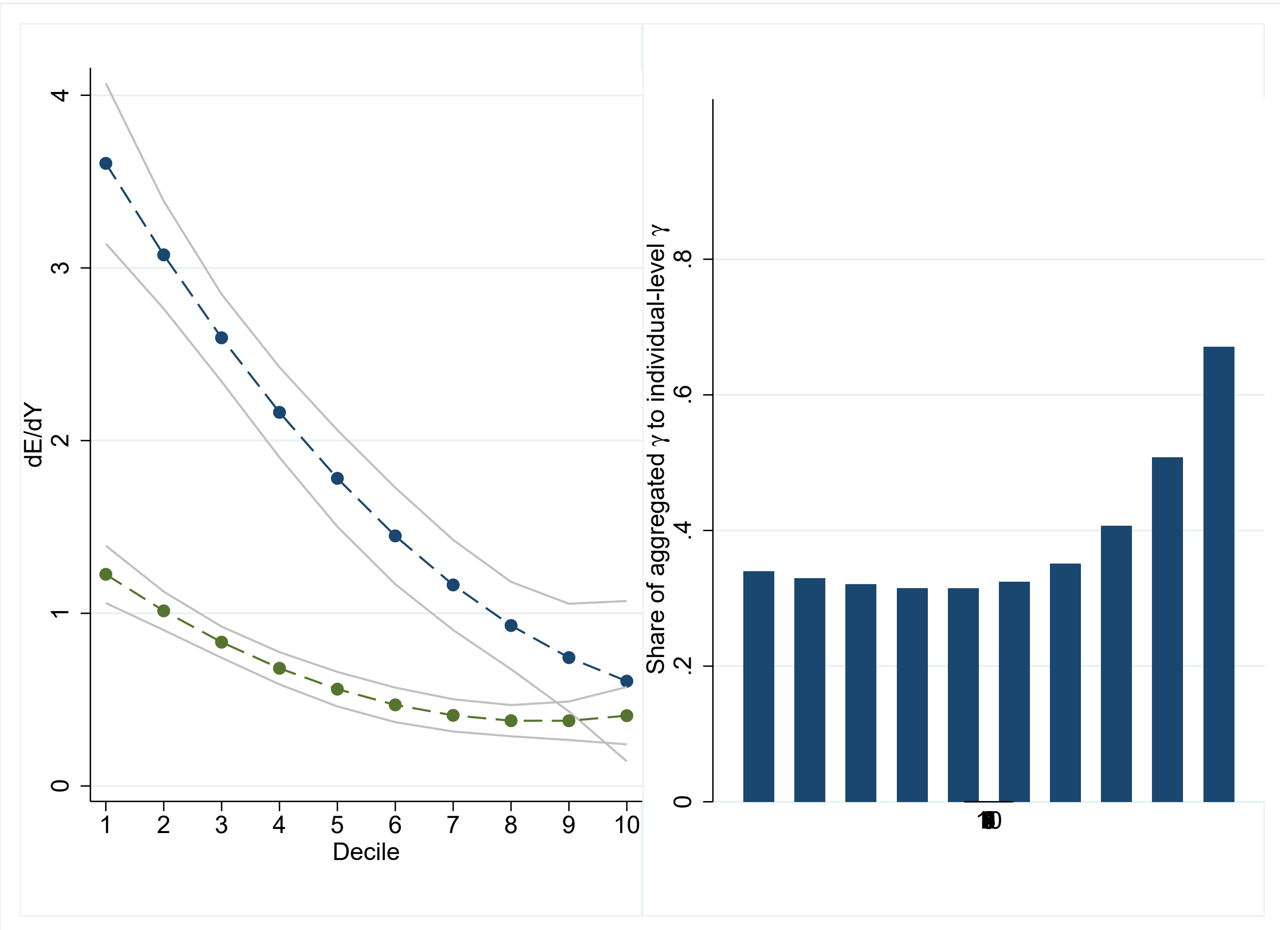

Figure 8: Left hand side: Earnings elasticity dE/dY by income decile estimated with individual and group-level (aggregated) HFCS data. Right-hand side: Ratio of the earnings elasticity at the group level to the elasticities estimated using the individual level data.

V. The Matching Multiplier

Next I turn to the question how the above evidence on heterogeneous earnings elasticities by households affects macroeconomic outcomes and macroeconomic policy. Earning shocks affect aggregate outcomes through many channels, but here I focus on the consumption response. The earnings heterogeneity of households will matter for macroeconomic outcomes, if there is a correlation between the earnings elasticity of the household type and its MPC. The relationship between the earnings elasticity and the MPC is a key object for models with households heterogeneity, as shown by Bilbiie (2020).

To understand this relation better we can decompose the standard definition of the MPC to quantify the covariance in this expression of the aggregate MPC in the economy:

$$MPC = \sum \frac{dC_h}{dE_h}\frac{dE_h}{dY} = \sum \frac{dE_h}{Y}\frac{dC_h}{dE_h} + \left(Cov\frac{dC_h}{dE_h}, \gamma_h \right) \tag{8}$$

where C_h is the consumption of household h, Y is aggregate output, and γ_h is the elasticity of household earnings to aggregate output. The MPC of an economy can be calculated using the first term – it is the sum of consumption responses of households to the aggregate income change. Yet implicitly behind the response of the macro aggregates, there is a more profound process taking place at the micro level. The aggregate output and therefore aggregate income change dY leads to changes in individual household earnings dE_h. Individual household consumption then responds to the individual household earnings change. The aggregate MPC can be presented in a second way – the aggregate MPC is given by the earnings weighted sum of all the individual MPCs – the sum of the consumption response of each household to their individual earnings change plus the covariance between the household MPC and the change in the individual earnings to aggregate income. The second term in this covariance is the earnings elasticity of individual earnings to aggregate output changes, that we have estimated already above - $\gamma_h = \frac{dE_i}{dY}\frac{Y}{E_i}$. While Bilbiie (2008) and Werning (2015) analyse theoretically the implications of different values for $\gamma_h = \frac{dE_i}{dY}\frac{Y}{E_i}$, they do not provide empirical evaluation of these different values.

Thinking about the implications of this decomposition, we should think about two possible cases. If the covariance is 0, then the heterogeneity of earnings elasticities does not matter – the earnings change of all workers change in the same proportion to the aggregate income output change. If however the covariance is positive - higher than 0, this means that it is exactly those workers that have higher MPCs, which also have higher earnings elasticities and are therefore more exposed to aggregate output changes.

By constructing the dataset of earnings changes from the HFCS, I provide the first dataset for the Euro area with panel data on household income changes, as well as with MPCs reported by households from the same source. The third and fourth wave of the HFCS provides data on MPCs by including a unique question directly asking households about their consumption response to an unexpected income shock. Survey participants are asked the question:

"Imagine you unexpectedly receive money from a lottery, equal to the amount of income your household receives in a month. What percent would you spend over the next 12 months on goods and services, as opposed to any amount you would save for later or use to repay loans?"

These self-reported MPCs are only available in the third and the fourth wave of the HFCS. By introducing a hypothetical question on households' marginal propensity to consume (MPC), we obtain self-reported estimates for the annual MPC. The question, framed in terms of an unexpected lottery winning, addresses the challenge of identifying unanticipated income shocks without assuming a structural relationship between consumption and income. This approach yields MPC point estimates for all households, facilitating the analysis of influencing factors, and therefore presents us with the full distribution of MPCs. The question encompasses overall household spending, crucial for studying the effects of redistributionary channels of monetary policy on aggregate spending. These self-reported MPCs arguably suffer some limitations, including bunching of results at round numbers, and potential asymmetry in responses to positive and negative income shocks. Despite these shortcomings, the HFCS results offer the first cross-country comparable estimates of MPCs for the Euro area. Combined with my estimated earnings elasticities, they can be used to quantify the effects of the relationship between the two objects, essential for any macroeconomic dynamics.

To get a basic understanding of the relationship between earnings elasticities and MPCs, I first compute a generalised figure to show how the two objects relate to each other in my dataset. The scatterplot in Figure 9 reports the correlation between the average earnings elasticities for each of the ten deciles, estimated in my HFCS dataset above, and average current MPC of this income decile. A clear positive relationship can be observed between the household betas and the mean MPCs by income decile, calculated using the self-reported MPCs in the fourth wave of the HFCS.

Next, I pursue to estimate the relationship between raw reported MPCs from the survey and the estimated earnings elasticities econometrically. I want to estimate the covariance between earnings elasticity and the MPC of different groups – the object $\left(Cov\frac{dC_h}{dE_h}, \gamma_h\right)$. This covariance will then enable me to quantify the degree to which the earnings of households with different MPCs are differently sensitive to the aggregate GDP change.

Patterson (2023) is the first to explore quantitatively this channel by using administrative US data and reports a so-called matching multiplier mechanism - that workers with high MPCs often have jobs which are highly impacted by downturns. Patterson (2023) estimates the covariance between household betas and households MPCs and finds that the mechanism is quite strong in the US and contributes to the amplification of recessions. The approach proposed also enables me to study how the fiscal multiplier is affected when the earnings elasticities of different households' income group vary systematically with the MPC. Following Patterson (2023), I estimate in the following a baseline estimation given by:

$$\Delta logE_{h,t} = \beta_0 + \beta_1 MPC_h + \beta_2 MPC_h \times \Delta logY_t + \epsilon_{h,t} \tag{9}$$

Where $\Delta logE_{h,t}$ is the 3-year change of income of household h, MPC is the self-reported household response to the MPC question in the survey, and $\Delta logY_t$ is the 3-years change in GDP. The coefficient β_2 is the coefficient of interest for us. It enables us to then evaluate how the amplification of shocks through the MPC is affected by the unequal incidence of aggregate shocks.

I therefore estimate the above equation using the HFCS data. I take the following data - $\Delta logE_{h,t}$ is the multi-year cumulated change of income of household h between two waves of the HFCS, as above; MPC is the self-reported household response to the MPC question of the same household in the previous wave of the survey, and $\Delta logY_t$ is the multi-year change in GDP, accumulated over the same period as for earnings. I take both the earnings changes and aggregate output changes therefore in the same way as above when evaluating household betas.

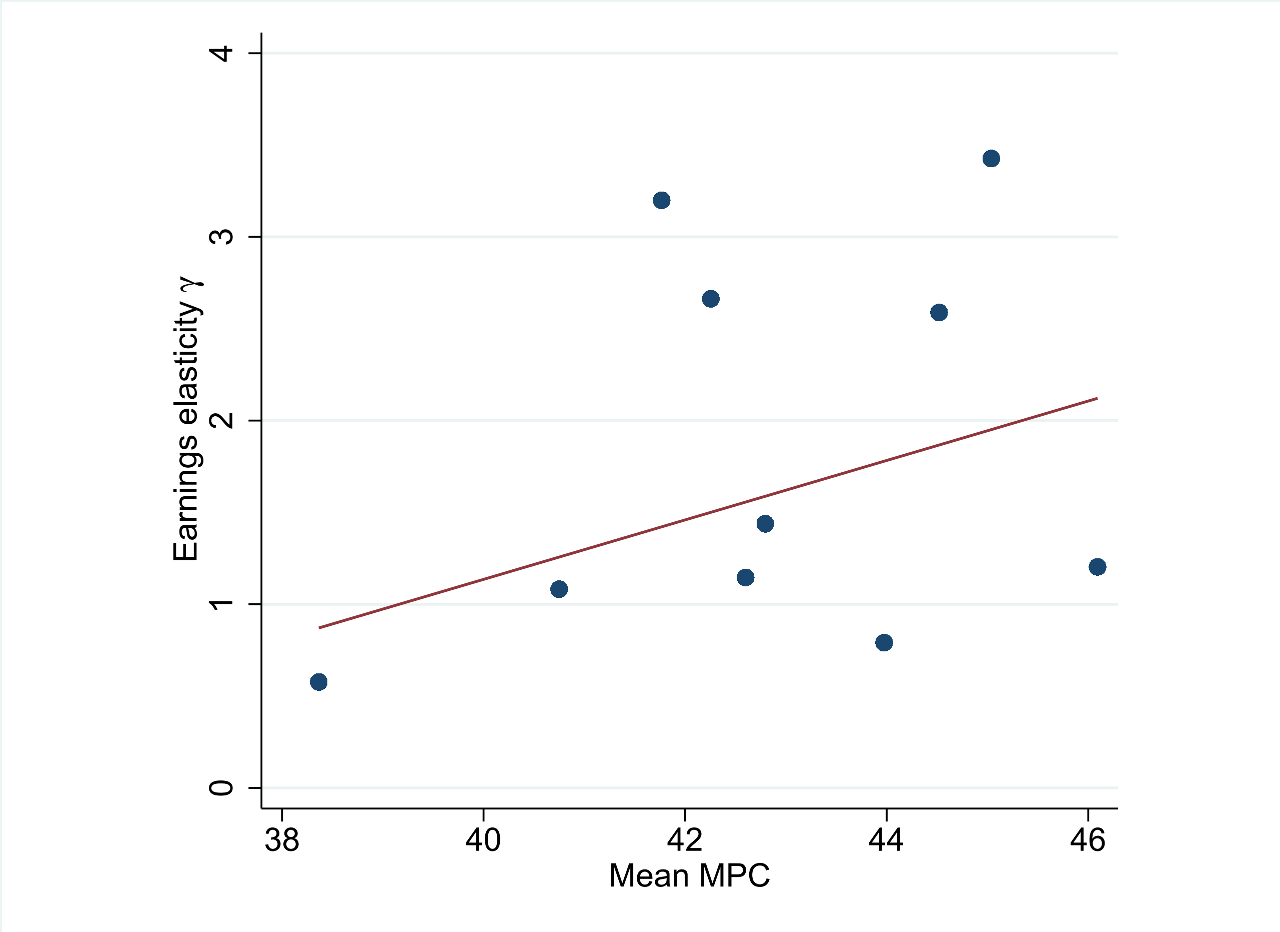

Figure 9: Relationship between earnings elasticities, estimated with the HFCS dataset, and mean MPCs of income deciles using self-reported MPCs in the HFCS, Euro area. : Graph reports the average earnings elasticity for each decile estimated with the HFCS individual data and the average MPC of this bin current liquid asset decile of households and their contemporary reported MPCs.

Table 1 summarizes the results of my baseline regression and reports the coefficients of interest, average MPC and the covariance between MPCs and earnings elasticities γ. The covariance between the MPC of the household and their estimated earnings elasticity γ is calculates as the product of the estimated coefficient β_2 from Equation 9 and the variance of the MPC. I run the regression for two specifications - for the full sample of MPC values and for a trimmed sample. The trimmed sample excludes households with a reported MPC of 0 and 1. While these are the most often reported values in the sample, they can bias my estimation downwards as the interaction term in the regression will be biased towards 0. In the regression with trimmed MPCs therefore I keep only households with reported MPCs above 0 and below 1. The resulting covariances is 0.065 for the full sample and 0.128 for the trimmed sample.

Based on these results, I can now also estimate the significance of the extra amplification mechanism

Table 1: Covariance between earnings elasticities and MPCs

	Baseline	Trimmed MPCs
Coefficient on $\mathrm{MPC_h} \times \Delta \mathrm{logY_t}$	0.487	3.813
Average MPC	0.445	0.456
Var(MPC)	0.134	0.033
Cov$(\mathrm{MPC_i}, \gamma_i)$	**0.065**	**0.128**

Table 2: Estimated aggregate MPCs with and without earnings elasticity based on calculated covariances

	Baseline	Trimmed MPCs
Benchmark MPC	0.250	0.250
MPC with gamma heterogeneity	0.294	0.336
Increase in MPC in %	**17.5**%	**34.2**%
Benchmark Multiplier	1.33	1.33
Multiplier with gamma heterogeneity	1.42	1.51
Increase in Multiplier %	**6.7**%	**13.5**%

due to the additional relationship between earnings elasticities and MPCs. The aggregate multiplier from an aggregate shock to the economy may be amplified in the case that agents in the economy have different income elasticity to GDP changes (in comparison to the case of the same income elasticity) if their consumption response also systematically varies with these changing elasticities. To estimate the magnitude of these two cases, I follow Patterson (2023) and define a benchmark, classical aggregate MPC and an amended MPC, where the amplification due to earnings elasticity heterogeneity is taken into account. The two different aggregate MPCs can be expressed as the following:

$$MPC_b = \alpha_l \left(\bar{\gamma}\overline{MPC}) + (1 - \alpha_l\bar{\gamma}\right) MPC_{nl}(10)$$

$$MPC_a = \alpha_l \left(\bar{\gamma}\overline{MPC} + \text{Cov}\left(MPC_i, \gamma_i\right)\right) + (1 - \alpha_l\bar{\gamma}) MPC_{nl} \tag{11}$$

To compute these values, I need information on α_l - the share of wages in total income in the economy; the MPC_{nl} – an assumption on the average MPC out of non-labor income, and $\bar{\gamma}$ - the elasticity of aggregate income to aggregate GDP in the economy, which is the measure of the earnings elasticity of workers to GDP changes when the elasticity is homogenous. I assume a share of wages in the economy of 2/3 and an MPC for non-labor income of 0.16, as assumed by Patterson (2023). I estimate the elasticity of aggregate income to aggregate GDP. For aggregate income I use data on wages and income from Eurostat. For the period 1995-2022 for the Euro area I estimate α_l =0.46. With these values I therefore compute the two aggregate MPCs. I also derive multipliers, calculated as in a standard macroeconomic textbook - one traditional, benchmark multiplier from the benchmark MPC value and another one from the amplified MPC value with earnings elasticity heterogeneity. The results are reported in Table 2 and in Figure 10.

The benchmark MPC for my baseline case equals 0.247, whereas the one with earnings elasticity equals 0.285. In the case of trimmed MPCs, which is econometrically more robust, the MPC increases further to 0.336. This is a considerable increase in the aggregate MPC between the two economies of 17.5% and 34.2% respectively. These aggregate MPCs can then be transformed into aggregate multipliers using the standard fiscal multiplier formula. The benchmark MPC of 0.250 results in an aggregate multiplier of 1.33. The two MPCs with γ heterogeneity result in an aggregate multiplier of 1.42 and 1.51 respectively. The heterogeneity of earnings elasticities and the relationship to heterogeneity in MPCs thus results in an increase in standard fiscal multipliers by between 6.7% and 13.5% in comparison to the case without inequality in earnings elasticities.

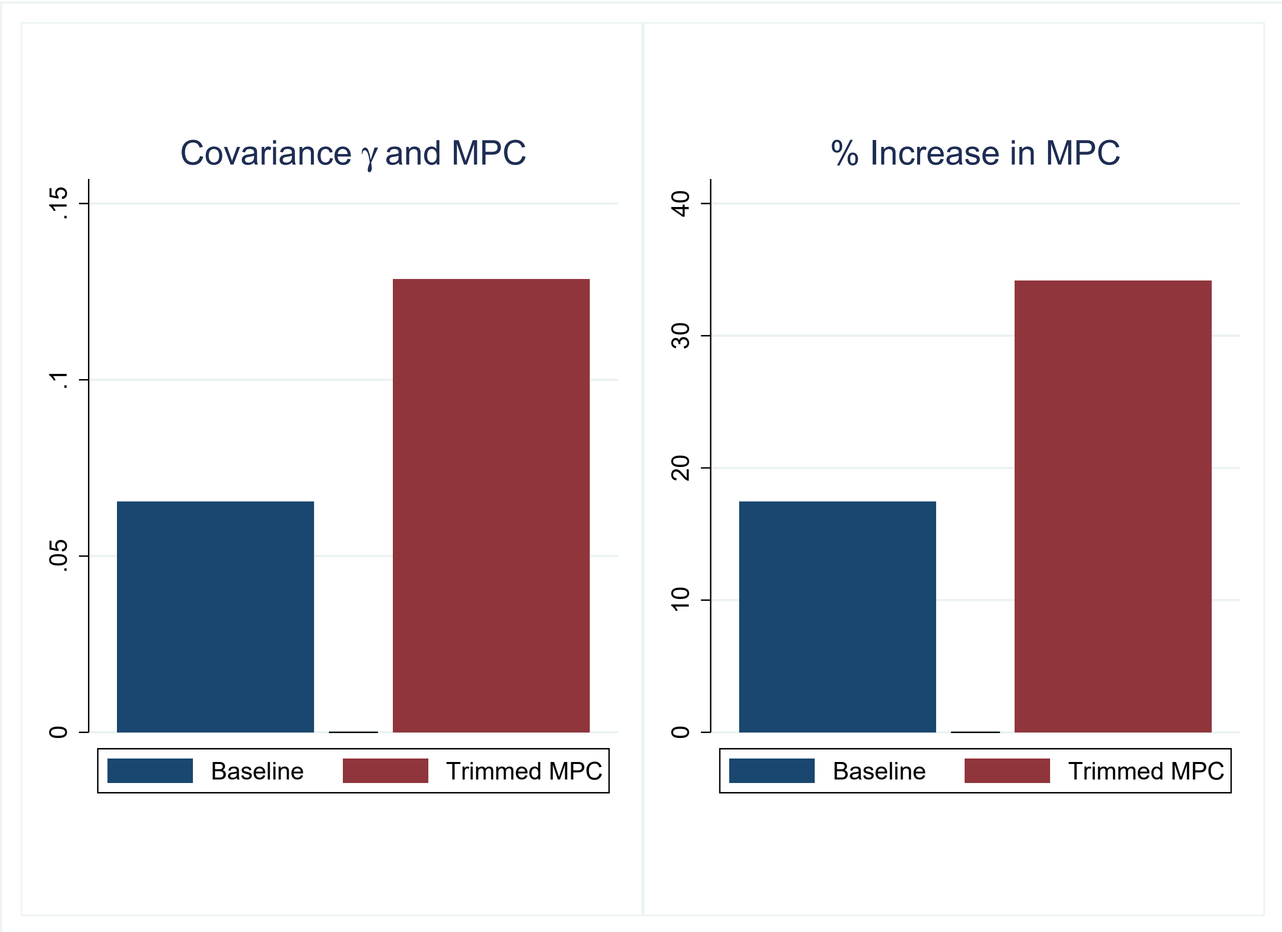

Figure 10: Covariance between earnings elasticities and MPCs and change in aggregate MPCs.

VI. MPCs and Liquid Asset Holdings

The theoretical models developed by Auclert (2019) and Patterson (2023) identify a way to use the relation between MPCs and income changes across the income distribution as a key moment for models with macroeconomic heterogeneity. At the same time however, the literature on MPC heterogeneity finds as the most important factor explaining the heterogeneity of MPCs across households not their position in the income distribution, but rather in the distribution of liquid wealth (Japelli Pistaferri 2014; Carroll, Slacalek, Tokuoka White 2017). In this section I analyse also this relationship between MPCs and liquid asset holdings.

Figure 11 plots earnings elasticities by deciles of the distribution of liquid assets. The difference between our previous estimations is that I now group the ten deciles for which I estimate the earnings beta by liquid asset holdings, as reported in the HFCS derived variable for net liquid assets. Figure 12 then groups the estimated earnings elasticities and plots them against the mean MPCs at each decile of the liquid assets distribution. The estimation for earnings elasticities delivers a mixed picture, where the clear downward pattern of decreasing earnings elasticities along the distribution does not hold. Interestingly however, the positive pattern between mean MPCs per decile and the earnings elasticity of different deciles, presented in Figure 12, still shows a positive pattern. This is especially the case for the contemporary liquid assets of households and their MPCs, shown in the right hand-side of the figure.

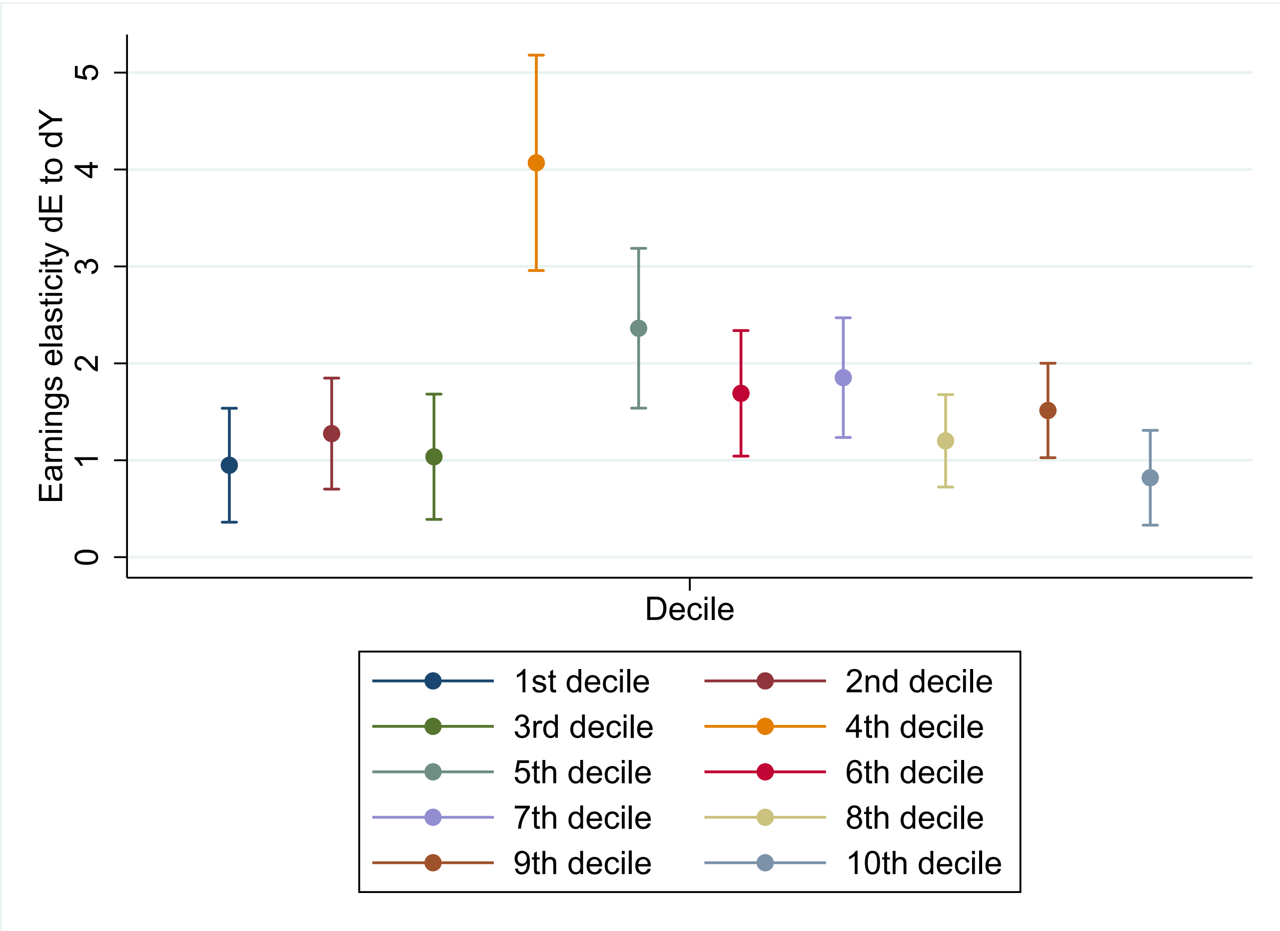

Figure 11: Earnings elasticity dE/dY using the HFCS dataset for the Euro area by liquid assets
Note: Earnings elasticities by decile of liquid assets using after-tax household earnings and excluding extreme observations where the rate of change was above 1000 percent (more than 10-fold increase or decrease in income).

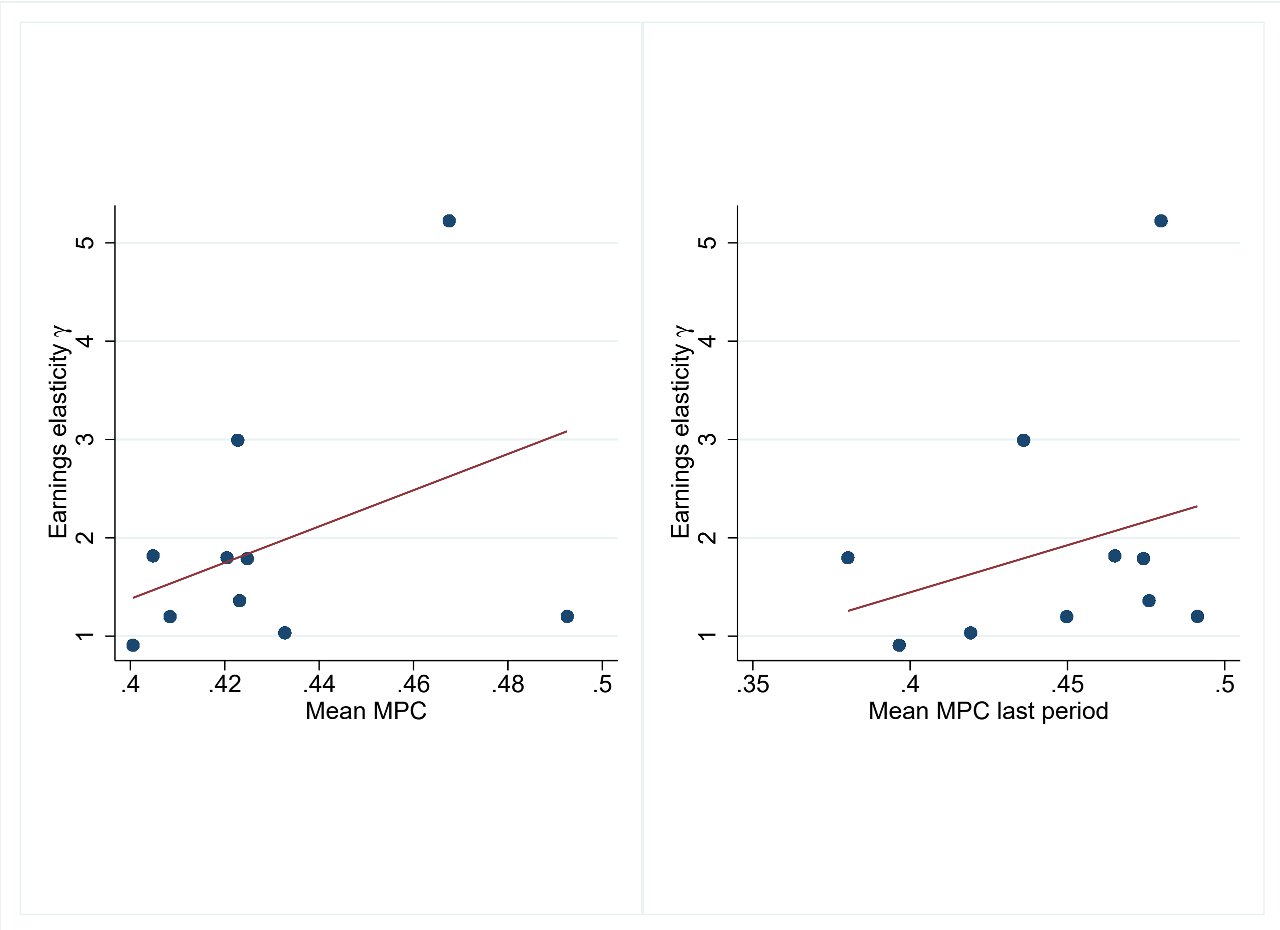

Figure 12: Relationship between earnings elasticities by deciles of liquid assets and mean MPCs of liquid wealth deciles using self-reported MPCs in the HFCS, Euro area.
Note: Left-hand side graph uses current liquid asset decile of households and their contemporary reported MPCs. Right-hand size uses the MPCs of households from the previous wave, following Patterson (2023). Patterson (2023) uses imputed MPC estimates when directly reported ones are not available and bases them on the previous period characteristics of the household, whereas I can directly use them.

VII. Conclusions

I document how the earnings of different household groups vary with the variation of aggregate output using a sample of European economies. All three datasets I use deliver a similar pattern – the incomes of the lowest income groups are much more sensitive to aggregate output changes than those in the middle of the income distribution. There is a clear downward pattern of the sensitivity of individual income to aggregate output changes - which we call household earnings betas, and these patterns hold up until the very high income bins, where the sensitivity again increases for the highest earning groups.

This is the first paper to obtain these results for the Euro area comparing both aggregated data by income groups and data compiled from panel household survey data. The results from the two types of sources are similar and also in line with previous estimates for the US, reported by Guvenen et. al (2017). The results estimated with the HFCS panel data point to higher earnings elasticities for the bottom income deciles than those from administrative data, as the panel data is best suited to also take into account income mobility across deciles. It thus enables the measurement of income changes for households which move along the income ladder, unlike data reported in aggregated form as in the other two datasets I use. A downward sloping pattern of earnings elasticities holds across all three datasets, whereas the difference is in the magnitude of these household betas.

Using the HFCS I also show the difference between using individual level data and aggregated data reported per income group or decile. I compare earnings elasticities using the same data and both approaches. At the left end of the distribution the individual level estimates give us extra information and deliver much higher elasticities to aggregate output fluctuations - because individual level data unlike aggregate data does not hide income mobility and income increases across bins it is a much better measure of income elasticities.

I also analyse how the correlation between earnings elasticities of different income groups and their MPCs affect macroeconomic dynamics. My results show that the groups with the highest earnings elasticities on average also have higher MPCs – a finding in line with Patterson (2023) for the US. In my baseline estimation, the covariance between earning elasticities of different income groups and their MPCs is between 6.5% and 12.8%. This finding has important repercussions for the reaction of consumption to aggregate changes – households with highest sensitivity to economic changes also have the highest MPCs, which acts as an amplification mechanism to any shock or policy change. To evaluate this, I calculate the size of the aggregate MPCs with and without this amplification mechanism. Because of the tight relation between heterogeneity of earnings elasticities and heterogeneity in MPCs, this amplification mechanism increases the aggregate multiplier of the economy by between 18% and 24%.

References

[1] A.K. Bartscher, M. Kuhn, M. Schularick, and P. Wachtel. "Monetary Policy and Racial Inequality - FEDERAL RESERVE BANK of NEW YORK". In: *Federal Reserve Bank of New York Staff Reports* January 2021.959 (Jan. 2021).

[2] Christian Bayer, Volker Tjaden, Lien Pham-Do, and Lütticke Ralph. "Precautionary Savings, Illiquid Assets, and the Aggregate Consequences of Shocks to Household Income Risk". In: *CEPR Discussion Papers* 10849 (Sept. 2015). URL: https://cepr.org/publications/DP10849 (visited on 11/16/2023).

[3] F.O. Bilbiie. "Limited Asset Markets Participation, Monetary Policy, and (Inverted) Aggregate Demand Logic". In: *Journal of Economic Theory* 140 (2008), pp. 162–196.

[4] F.O. Bilbiie. "Monetary policy and heterogeneity: An analytical framework". In: *CEPR Discussion Paper* DP12601 (2018). URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3106805.

[5] A. Castañeda, J. Dıaz-Giménez, and J.-V. Rıos-Rull. "Exploring the income distribution business cycle dynamics". In: *Journal of Monetary Economics* 42 (1998), pp. 93–130.

[6] M. Dossche, J. Slacálek, and G. Wolswijk. "Monetary policy and inequality". In: *ECB Economic Bulletin* 2/2021 (2021). URL: https://www.ecb.europa.eu/pub/economic-bulletin/articles/2021/html/ecb.ebart202102_01~1773181511.en.html.

[7] Maarten Dossche and Jacob Hartwig. "Household income risk over the business cycle". In: *Economic Bulletin Boxes* 6.5 (2019).

[8] European Central Bank. "Household income risk over the business cycle". In: *Economic Bulletin* 6 (2019).

[9] F. Guvenen, S. Schulhofer-Wohl, J. Song, and M. Yogo. "Worker Betas: Five Facts about Systematic Earnings Risk". In: *American Economic Review: Papers and Proceedings* 107.5 (2017), pp. 398–403.

[10] Marcus Hagedorn, Iourii Manovskii, and Kurt Mitman. "The Fiscal Multiplier". In: *NBER Working Paper* 25571 (Feb. 2019), w25571. DOI: 10.3386/w25571. URL: http://www.nber.org/papers/w25571.pdf (visited on 11/16/2023).

[11] J. Heathcote and F. Perri. "Wealth and Volatility". In: *The Review of Economic Studies* 85.4 (Oct. 2018), pp. 2173–2213. DOI: 10.1093/restud/rdx074.. URL: https://doi.org/10.1093/restud/rdx074..

[12] Tullio Jappelli and Luigi Pistaferri. "Fiscal Policy and MPC Heterogeneity". In: *American Economic Journal: Macroeconom-ics* 6.4 (2014), pp. 107–36.

[13] Tullio Jappelli and Luigi Pistaferri. "The Consumption Response to Income Changes". In: *Annual Review of Economics, Annu-al Reviews* 2.1 (2010), pp. 479–506, 09.

[14] G. Kaplan and G.L. Violante. "Microeconomic Heterogeneity and Macroeconomic Shocks". In: *Journal of Economic Perspectives* 32.3 (Aug. 2018), pp. 167–94. DOI: 10.1257/jep.32.3.167.. URL: https://doi.org/10.1257/jep.32.3.167..

[15] Greg Kaplan, Benjamin Moll, and Giovanni L. Violante. "Monetary Policy According to HANK". In: *American Economic Re-view* 108.3 (2018), pp. 697–743.

[16] John Maynard Keynes. *The General Theory of Employment.* Place: Book III. Macmillon and Co., 1936.

[17] A. McKay. "Time-varying idiosyncratic risk and aggregate consumption dynamics". In: *Journal of Monetary Economics* 88 (2017), pp. 1–14.

[18] Jonathan A. Parker and Annette Vissing-Jorgensen. "Who Bears Aggregate Fluctuations and How?" In: *American Economic Review* 99.2 (2009), pp. 399–405.

[19] C. Patterson. "The Matching Multiplier and the Amplification of Recessions"". In: *Society for Economic Dynamics, Meeting Pa-pers* 95 (2019).

[20] Iván Werning. "Incomplete Markets and Aggregate Demand". In: *NBER Working Paper* 21448 (Aug. 2015). DOI: 10.3386/w21448. URL: http://www.nber.org/papers/w21448.pdf (visited on 11/16/2023).

VIII. Annex

1. Countries Covered in the HFCS Panel Component and First Year of Panel

Table 3: Panel countries

Country	Wave 1	Wave 2	Wave 3	Wave 4
Belgium		X		
Germany		X		
Estonia			X	
Ireland				X
Spain		X		
France			X	
Italy		X		
Cyprus		X		
Lithuania				X
Malta		X		
Netherlands		X		
Slovakia			X	
Finland			X	