The joint distribution of wealth, income and consumption in Austria: a cautionary note on heterogeneity

Peter Lindner, Martin Schürz[1]
Refereed by: Juha Honkkila, European Central Bank

In this paper we analyze the joint distribution of wealth, income and consumption in Austria. We use data from three distinct surveys, each of which focuses on one of these components, and combine these data based on a statistical matching procedure. We find that statistical data matching does not overcome the problems connected with each of the underlying data sources but rather multiplies them. There is a likely tendency to the mean in the statistical matching procedure. Thus, the tails of the distribution emerge as particularly problematic. We document the enormous difference between the three indicators usually used for describing the joint distribution. These differences can be identified in particular for specific subgroups. Finally, we argue for using only one comprehensive source of data – the Household Finance and Consumption Survey (HFCS) – to estimate the joint distribution.

JEL classification: D30, I31, G51
Keywords: joint distribution, net wealth, statistical matching, HFCS

In this paper we address conceptual issues of measuring social inequality. We study the joint distribution of wealth, income and consumption in Austria. To this end, we utilize statistical data matching techniques that focus on the components of this joint distribution at the household level. In the empirical analysis, we use the data of three waves of the Household Finance and Consumption Survey (HFCS) in Austria and the European Statistics on Income and Living Conditions (EU-SILC) as well as two waves of the Households Budget Survey (HBS).

A key contribution of our paper is to highlight problems of experimental statistics. The relationship between wealth, income and consumption does not get clearer by matching data from different surveys. Issues that remain unresolved are normative issues of justification of the weight of these variables in a joint distribution plus the specific determination of technical details. Statistical data matching might not create a useful joint distribution because it takes data from multiple surveys that have different strengths and weaknesses. In consequence, matching does not bring diverse data closer together but rather multiplies the data problems of different surveys. Thus, we underline the advantages of the HFCS data despite the fact that it covers the consumption variable only by a few questions.

This paper is structured as follows: Section 1 provides theoretical considerations and discusses the three variables considered in the analysis of the joint distribution. In section 2, we introduce the survey data underlying the analysis. In section 3, we present the statistical data matching method. In section 4, we provide our results, where we first analyze the differences in the composition of the variables for the joint distribution and its development over time. Then we look at distinct groups of households and include a link to the analysis of household vulnerability. The last

[1] *Oesterreichische Nationalbank, Economic Analysis Division, peter.lindner@oenb.at, martin.schuerz@oenb.at. Opinions expressed by the authors of studies do not necessarily reflect the official viewpoint of the OeNB or the Eurosystem. The authors would like to thank the referee and Pirmin Fessler (OeNB) for helpful comments and valuable suggestions.*

part of section 4 analyzes the tails of the joint distribution and their connection to the statistical matching process. Section 5 concludes.

1 Theoretical considerations

Former IMF Managing Director Christine Lagarde once claimed that "reducing excessive inequality [...] is not just morally and politically correct but it is good economics".[2] And on the occasion of the launch of a report on inequality, the Secretary-General of the OECD, Angel Gurría, declared that "inequality can no longer be treated as an afterthought. We need to focus the debate on how the benefits of growth are distributed."[3]

When looking into social inequality, should wealth, income and consumption be analyzed together or separately? Intuitively, we would expect a strong relationship between these variables. However, at least in developed countries, the Gini coefficient for household wealth is in the range between 0.65 and 0.85, and for post-tax income it is between 0.25 and 0.35. So far, most studies investigating household well-being take a partial view by looking at wealth, income or consumption separately. But since the publication of the Stiglitz-Sen-Fitoussi Report on the measurement of economic well-being 2009, it has been clear that we should take a more comprehensive view of households.

Hence, international institutions like Eurostat and the OECD have acknowledged the need to broaden the view when analyzing households. The joint distribution, which takes all three components into account, is thus gaining prominence in international discussions.

The three-dimensional distribution is at the core of this paper, which will also focus on the characteristics of different groups of households in society. Do differences between household groups with respect to the joint distribution pose potentially different threats to financial stability? Using all three waves of the HFCS and the available information from the other two surveys, we can analyze a timespan of about seven years. Thus, we can get a first idea of how the composition of specific groups changed in the recent past. And how are the tails of the distribution affected if we take into account the multidimensionality as well as different variants of statistical matching? These are the questions that we analyze in this paper.

The annex to this study provides more detailed information about the matching technique that we applied to generate the single data source for our analysis. We discuss the explicit and implicit assumptions needed in the matching procedures and investigate how sensitive our results are with respect to these assumptions.

The conceptual challenges of a multidimensional approach in understanding wealth, income and consumption are considerable. Assigning weights within multidimensional inequality to single variables is a difficult task. "Equality of what" is the title of a seminal contribution by Nobel prize winner Amartya Sen. Sen's approach was operationalized in the Human Development Index produced by the UN Development Program, which focuses on life expectancy, education and income. As we look at wealth, income and consumption, we deviate from this approach.

[2] *Taken from a speech given in Brussels on June 17, 2015, available at https://www.imf.org/en/News/Articles/2015/09/28/04/53/sp061715.*

[3] *Taken from the launch of "In It Together – Why Less Inequality Benefits All" held in Paris on May 21, 2015, available at https://www.oecd.org/social/publication-launch-in-it-together-why-less-inequality-benefits-all.htm.*

But an all-encompassing metric is still not available and the heterogeneity in different forms of inequality is easily overseen. For instance, a person may have accumulated substantial wealth over their working life, but following retirement they have a modest level of income. On the other hand, young people may have high incomes but may have not yet had the opportunity to generate substantial assets. Theoretically, we should be concerned with inequality across the entire distribution. But the notions of (top) percentiles are not the same for the distribution of income, wealth and consumption. They can be correlated strongly or less strongly, and, in any case, they will be different across countries. Moreover, the classification of percentiles is mostly ad hoc and arbitrary in inequality studies. Whether one focuses on the bottom 50% or the bottom 10% will make a huge difference in substance.

As will be seen below, the distributions of wealth, income and consumption can hardly be compared directly. Piketty (2014) assumes that "the future structure of inequality might bring together extreme forms of domination based simultaneously on property and culture (in brief: Marx and Bourdieu reconciled)" (Piketty, 2014, p. 743). This broader view has not been conceptualized up to now. Groups defined along indicators that are wealth based – such as renters, owners, and capitalists – are not necessarily the same as groups defined along income and consumption indicators.

Many studies on inequality, particularly in developing economies, have focused on consumption or expenditure. And theoretically, people are expected to smooth consumption over their lifetime. Measuring consumption often gives a more direct estimate of well-being than income as particularly in developing economies, income is hard to measure because of non-market activity. Consumption possibilities are determined by currently earned income and accumulated wealth and by the possibility to borrow against an existing stock of wealth.

Income indicates the ability to meet material needs in the short term. Because of income taxation, long-time series on income are available in most countries. The World Inequality Database brings together estimates for numerous OECD countries (see: *https://wid.world/*). The median of disposable household income is an indicator of normal living standard. These statistics, however, abstract from the source of this income. Income from labor is inherently different from income from capital. The mechanisms that determine labor income include supply and demand for different skills, the state of the education system, institutions that affect the labor market and the determination of wages. For capital income, by contrast, savings and investment behavior play a decisive role.

Wealth measures the private ownership of assets. Wealth is more stable than income over time and less reliant on personal effort. Laws governing inheritances and gifts and the functioning of the real estate and financial markets matter a lot (see Piketty, 2014, p. 243; Pistor, 2019). There are numerous methodological challenges in studying wealth. Due to the sensitivity of the topic, measuring wealth involves participation and reporting problems to an even larger extent than income and consumption surveys. In particular, wealthy people are less likely to participate in voluntary surveys. Additionally, wealth has a fictitious component. Its true value is determined on the market only once it is liquidated. Furthermore, its composition is a topic of discussion, e.g. the inclusion of human capital, social capital or pension wealth.

Table 1

Schematic overview of underlying data and documentation

HFCS (basis, three waves every 3 years)	**EU-SILC** (yearly, from 2003 to 2018)	**HBS** (every 5 years, starting 1999/2000)
2010/**11**: Wave I (almost equal share of households in each year)	EU-SILC wave 2010	HBS 2009/10
Fessler et al. (2012) and Albacete et al. (2012)	Statistics Austria (2012)	Statistics Austria (2013)
20**14**/15: Wave II (almost all households interviewed in 2014)	EU-SILC wave 2014	HBS 2014/15
Fessler et al. (2016) and Albacete et al. (2016)	Statistics Austria (2016)	Statistics Austria (2018b)
2016/**17**: Wave III (almost all households interviewed in 2017)	EU-SILC wave 2017	HBS 2014/15
Fessler et al. (2018) and Albacete et al. (2018)	Statistics Austria (2018a)	Statistics Austria (2018b)

Source: Authors' compilation.

There are, however, subfields in the literature where the results from a multi-dimensional approach are combined. Household vulnerability and fragility based on micro data commonly take into account the three indicators wealth (net wealth as well as gross wealth), income and consumption. Debt-to-income (DTI) and debt service-to-income (DSTI) ratios, for example, are based both on asset components and income. In macroprudential policy, which sets limits on DTI and DSTI, consumption is taken on board, i.e. it is considered how much income a household needs for basic consumption in order to still be able to service its debt (see Albacete and Lindner, 2013, for an early attempt for Austria).

2 Data

There is no single data source in Austria that covers all aspects of wealth, income and consumption of households in sufficient detail. There are essentially three surveys that collect household micro data: The Household Finance and Consumption Survey (HFCS) conducted by the OeNB tackles the most difficult item, wealth, and Statistics Austria conducts two major surveys for households, the European Statistics on Income and Living Standards (EU-SILC) and the Household Budget Survey (HBS). The former mainly targets income while the latter has extensive information on consumption. As we want to shed some light on the development over time, we use multiple waves of each survey in this analysis.

Table 1 gives an overview of the data we use and their documentation. For the HFCS, the field period encompasses two calendar years, so we highlighted the year we take as the main reference year and state the share of households interviewed in each year.

We do not adjust the results for inflation over the waves. The main reason for that is that we are interested in the joint distributional information rather than the changes in levels. Applying a constant factor of inflation adjustment does not alter the ranking in a distribution but only shifts the level.

3 Matching basics

Analyzing a joint distribution of wealth, income and consumption would ideally require one data source covering all the necessary information. Since no survey in Austria covers all items for wealth, income and consumption, it is possible to use statistical matching techniques to come up with one dataset including all desired variables.[4]

In statistics, there are various approaches to matching data. D'Orazio et al. (2006) provide an overview, and the annex to this study ("Technical aspects of the matching process") discusses some matching possibilities. There are differences in the number of observations that remain missing after matching and the precision of the matching procedure. It is beyond the scope of this article to evaluate these pros and cons of different matching algorithms.[5] Instead, we simply opt for a stratified single random (rank) hotdeck procedure since it is commonly used in the literature and easy to implement. We analyze the sensitivity of this procedure with regard to the assumptions by experimenting with two distinct stratification implementations (see below). Essentially, this technique starts with one dataset. In our case, this is the HFCS, since information on wealth is hardest to impute. We look for similar households in a donor dataset, i.e. EU-SILC and the HBS. Once such similar households are found, one of the households from the donor data is randomly taken and the value of the desired information, i.e. income or consumption, respectively, attached to the specific household in the HFCS. Once this value is attached, all the survey specifics, such as weighting and imputation, are taken from the HFCS and into account for the estimation.

The stratification in this process defines how similar these matched households must be. It basically defines bins along dimensions of the sociodemographic information households are assigned to. Theoretically, for the statistical matching process to yield unbiased results, the stratification has to ensure that the conditional independence assumption (CIA) holds. This assumption says that given the bin, no other information can provide insight into the target variable. We use two different stratification definitions in order to show the impact on one of our central results. In the main matching ("matching I") procedure we use the following information for stratification in matching EU-SILC data to HFCS data:

- *Household type:*
  "single;" "couple (no children);" "more than two adults, no children;" "single with at least one child;" "couple with at least one child;" "more than two adults with at least one child"
- *Tenure status (household main residence):*
  "owner;"[6] "renter"
- *Floor space of the household's main residence:*
  Five quintiles in each of the underlying data sets[7]
- *Age of the household head:*[8]
  eight categories in ten-year steps
- *Education of the household head:*
  "without secondary;" "secondary education;" "tertiary education"[9]

[4] *The HFCS is the most complete data set but lacks some information on income, while the data from EU-SILC can draw from register information and has more information about the details of consumption.*

[5] *D'Orazio et al. (2006) are a good starting point for this discussion.*

[6] *Including households who live in their main residence free of charge.*

[7] *This basically assumes similarity of households in the same rank of the distribution of floor space.*

[8] *Household head is defined as the main income earner.*

[9] *The Austrian education system has various tracks to reach the highest level of education. Austrian education levels are then coded to international standards.*

The choice of the stratification variables is essential in this exercise. We made sure to include information on household structure, indicators for wealth (tenure status and floor space) and sociodemographic information. The inclusion of education is of particular importance due to its correlation with income. In the matching procedure of HFCS and HBS data, we use the same information as above except the age categories, which are reduced to three categories, and additionally use five quintiles of disposable household income.

For the less granular "matching II" of EU-SILC and HFCS data, we exclude the tenure status as well as the information on floor space and reduce the number of household types and age categories. Similarly, for the matching II of HBS and HFCS data, we exclude the tenure status as well as the information on floor space and reduce household type categories. We implement a single imputation of the matched households with no special attention paid to the additional statistical uncertainty.

4 Results

Analyzing the joint distribution of wealth, income and consumption, we focus first on a general overview and on the development over time. We then divide society into groups of households along a natural split emerging from the wealth distribution. Rounding up the empirical part of the study, we take a closer look at the tails of the distribution.

4.1 Overview

Table 2 allows us to compare the information from the three data sources. The results from EU-SILC and the HBS are those of the matched data. We display mean, median and the 5th and the 95th percentiles. Net wealth is reported as defined in the HFCS (see e.g. Fessler et al., 2018). For gross income, we show the equivalized yearly level, with the equivalization being based on the OECD method. The same is applied to monthly consumption. While consumption data from the HFCS are based on a single question that asks for total monthly consumption, the information from the HBS is a sum of all the various components of consumption.

Net wealth ranges from zero (5th percentile) to about EUR 900,000 (95th percentile), while yearly income ranges from EUR 12,000 to EUR 81,000, and monthly consumption ranges from below EUR 1,000 to above EUR 4,000. Wealth is much more skewed than income and consumption. This is the reason why income and consumption should be matched to information on wealth and not the other way round. Table 2 shows higher levels of income (except at the bottom of the distribution) and consumption in EU-SILC and HBS data compared to their HFCS counterparts. This is due to the focus of the respective surveys and the fact that the EU-SILC can access register income data.

The overview descriptive statistics, however, do not give any information about the joint distribution of the three

Table 2

Descriptive statistics

	P5	Median	Mean	P95
	EUR thousand			
HFCS				
Net wealth	0.0	82.7	250.3	865.9
Gross income	11.9	28.2	33.2	63.6
Consumption	0.4	0.7	0.8	1.3
EU-SILC				
Gross income	11.9	32.5	37.9	80.9
HBS				
Consumption	0.8	1.8	2.0	4.3

Source: HFCS 2017, OeNB; EU-SILC, HBS, Statistics Austria.

Chart 1

Joint distribution from 2010 to 2017

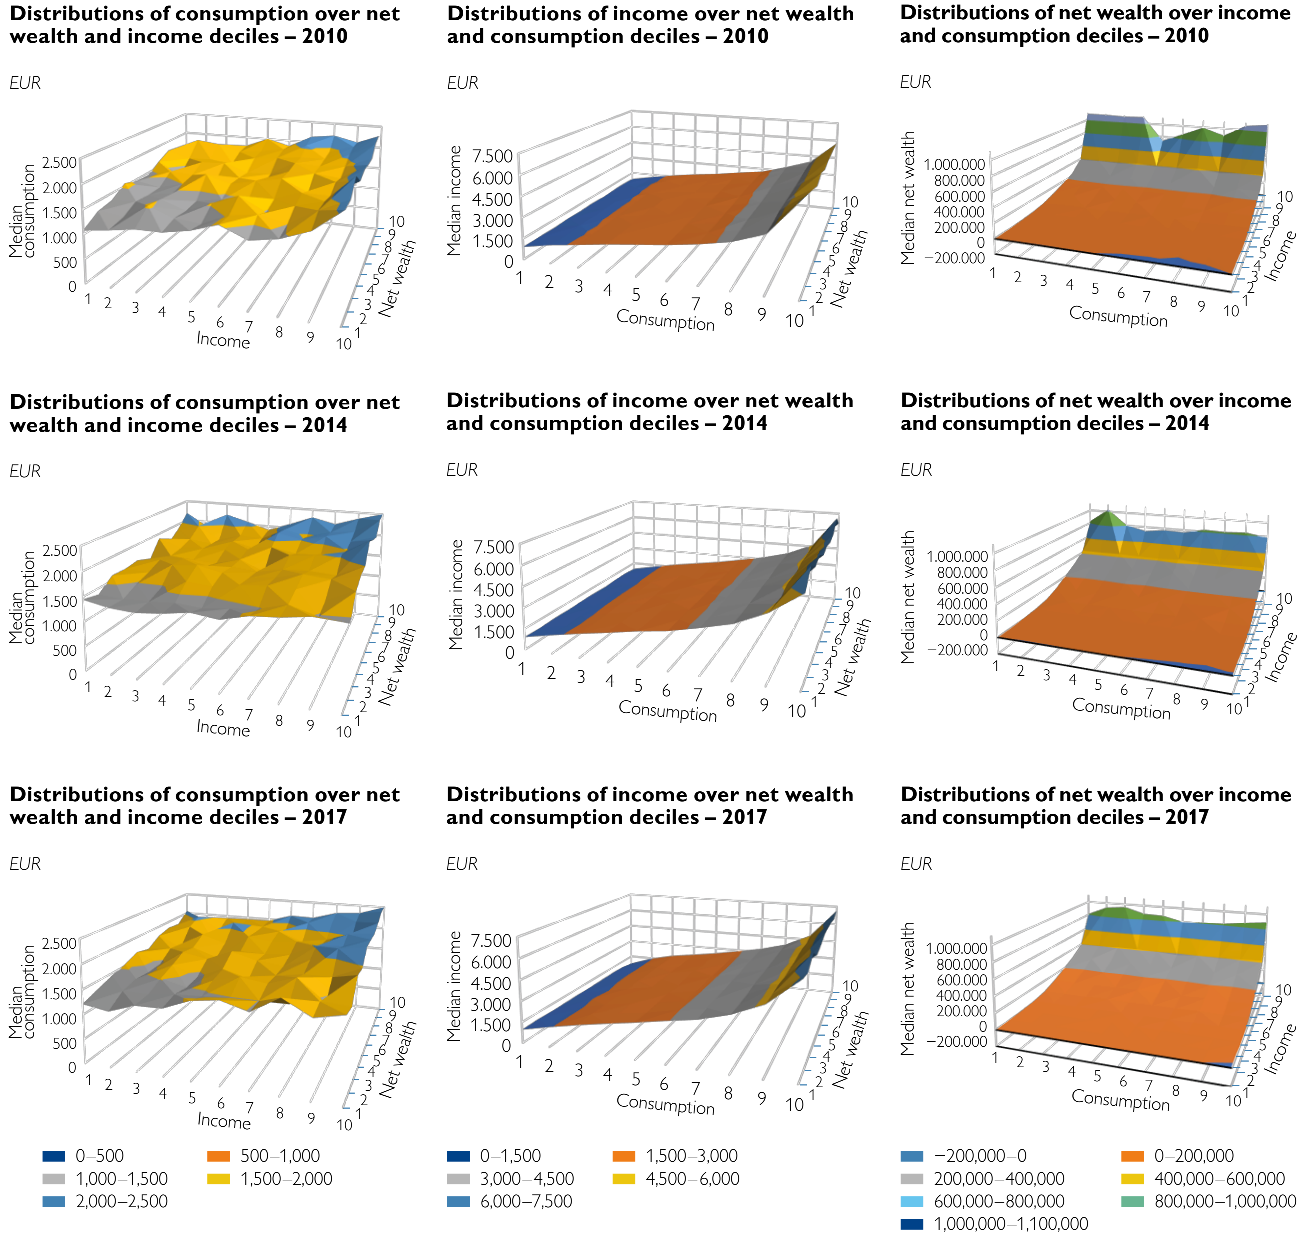

Source: HFCS 2010-2017, OeNB; EU-SILC 2010, 2014, 2017, HBS 2009/10 and 2014/15, Statistics Austria.

indicators together. To obtain such information, we split the sample of each of the combinations of two out of the three indicators according to ten deciles of equal population size in each variable. We sort the data according to one indicator – consumption, for instance – and calculate by using weights ten groups of society with increasing consumption. These groups are then crossed with the groups from the other indicator to obtain 100 subgroups of the whole population in each of the two-dimensional spaces, i.e. wealth–income, wealth-consumption, and income-consumption. As there are about 3,000 households in the sample, the number of households is in general large enough to estimate the results in a cell. Although it would be interesting to go beyond this split, data limitations in terms of small sample size are problematic.[10] For each of the cells we calculate the median of the third indicator in each subgroup. For example, we calculate deciles for net wealth and income. In each of the net wealth decile, there are households from each of the ten income deciles. The interaction of these two decile groups yields 100 cells. In each of the 100 cells, we calculate median consumption. This is shown in the bottom-left panel of chart 1 for the information in the year 2017.

We see that in general, consumption rises with income and wealth. There are, however, high levels of consumption also at low levels of wealth or at low levels of income. Due to high income, for example, high consumption can be afforded despite a lack of wealth. Also, for households with substantial wealth, a low level of income does not necessarily translate into low levels of consumption.

The panel in the middle at the bottom of chart 1 shows median income levels across the wealth-consumption space. The highest levels of income are found at the top of the consumption distribution. Also, the slope of income along the distribution of consumption becomes steeper, meaning that the level of income is relatively flat at the bottom and middle of the consumption distribution but rises sharply at the top of the consumption distribution. In other words, at the lower and middle consumption levels, income is relatively flat, while it increases sharply at the top. Over the wealth distribution, income is relatively flat, with the highest levels being reached at the bottom of the net wealth distribution combined with the highest level of consumption. These are households that display high levels of income which is then consumed. Looking at the data underlying chart 1, we see that median monthly income actually increases slightly with ascending net wealth deciles, so there is also a small positive slope along the distribution of net wealth.

Lastly, the bottom-right panel of chart 1 portrays the median level of net wealth in the income-consumption space. The level of net wealth is relatively flat along the consumption distribution while it increases along the income distribution dimension. Thus, there is a higher correlation between income and wealth than between consumption and wealth. This result hints at a satiation point to consumption for households. Even with high resources available there is a kind of upper boundary to consumption.

Additionally, we see that there are households with high levels of consumption and low levels of income (bottom-right corner in the bottom right panel of chart 1) that, accordingly, have very low, or even negative levels of net wealth. These are the households that are of special concern in the analysis of financial stability. It is also worth noting that, again, the range of values covered in the bottom-right panel of chart 1 dwarfs what we found for consumption and income in the two other bottom panels. This once again shows that focusing on wealth instead of income and consumption seems warranted (see e.g. Piketty, 2014).

In addition to the results (bottom panels) for the 2017 wave, chart 1 also shows the results for 2014 (middle row) and 2010 (top row). Overall, there seems to be stability over time. Each of the three-part sets displays a pattern similar to the one described above for 2017.

[10] *We repeated the whole exercise with 20 vingtiles for each indicator instead of ten deciles. The overall results remain unchanged. However, the crossing of two dimensions are resulting in estimations of medians for 400 cells which are sensitive to outliers.*

As the results are not inflation adjusted, the absolute level of consumption increased broadly over time. Furthermore, the highest levels of income increased noticeably while the remaining results in this regard over the consumption-net wealth space remained broadly stable.

4.2 Three groups of households aligned to joint distribution

In the following section, we look at the multidimensional distribution and distinguish the following three distinct groups of households (as done in Fessler and Schürz, 2018): households that rent their main residence (renter), households that own the main residence (owner) and households that own their main residence and/or have business wealth and/or have income from renting out real estate other than the household's main residence (capitalist). This split is done according to the data in the HFCS, meaning that the information necessary for the classification is taken from the HFCS.

Table 3 breaks indicators of the joint distribution down by these three groups.

Consumption remains roughly the same over the three groups. The average income of capitalists (about EUR 59,000) in the HFCS is twice as high than for renters (about EUR 29,000). Looking at the matched information from the EU-SILC, renters earn, on average an income of about EUR 45,000, which compares to EUR 44,000 for capitalists. This illustrates the issue of statistical matching and reversion to the mean along a dimension that is not part of the stratification process, such as the groups defined above. In terms of wealth, however, the three groups of households display completely different levels. Median net wealth for renters amounts to about EUR 15,000, while it is many times this amount for owners – EUR 250,000 – and even higher for capitalists – around EUR 650,000. Thus, levels of wealth differ substantially among the household groups, while levels are quite similar when it comes to income and consumption.

Over the whole distribution, we can analyze which type of households resides in which part of the distribution by looking at chart 2.[11]

Although the share of renters decreases over the income and consumption distribution and the share of owners increases therein, in both cases we cannot find the clear separation among groups of households known from the net wealth distribution. Furthermore, capitalists are spread out over both income and consumption distribution spaces.

Compared to their share in the overall population, both owners and capitalists have a disproportionately high share of wealth, income and consumption (see table A3). This result is stronger for capitalists than for owners. It is also a lot stronger for wealth than for income and consumption. The shares across

Table 3

Descriptive statistics over groups in society

	Renter		Owner		Capitalist	
	Median	Mean	Median	Mean	Median	Mean
	EUR thousand					
HFCS						
Net wealth	14.8	58.2	252.2	323.4	648.6	1,176.4
Gross income	25.6	28.7	30.6	34.0	38.1	58.8
Consumption	0.7	0.7	0.7	0.9	0.7	0.8
EU-SILC						
Gross income	39.3	45.0	29.1	32.8	36.9	43.7
HBS						
Consumption	2.0	2.4	2.0	2.3	1.6	1.8

Source: HFCS 2017, OeNB; EU-SILC, HBS, Statistics Austria.

[11] *These charts are based on the smoothed results of a kernel-weighted local polynomial regression with an Epanechnikov kernel function, a degree of 1 and a half width of 5.*

Chart 2

Share of groups of households

Share of groups of households across net wealth | **Share of groups of households across income** | **Share of groups of households across consumption**

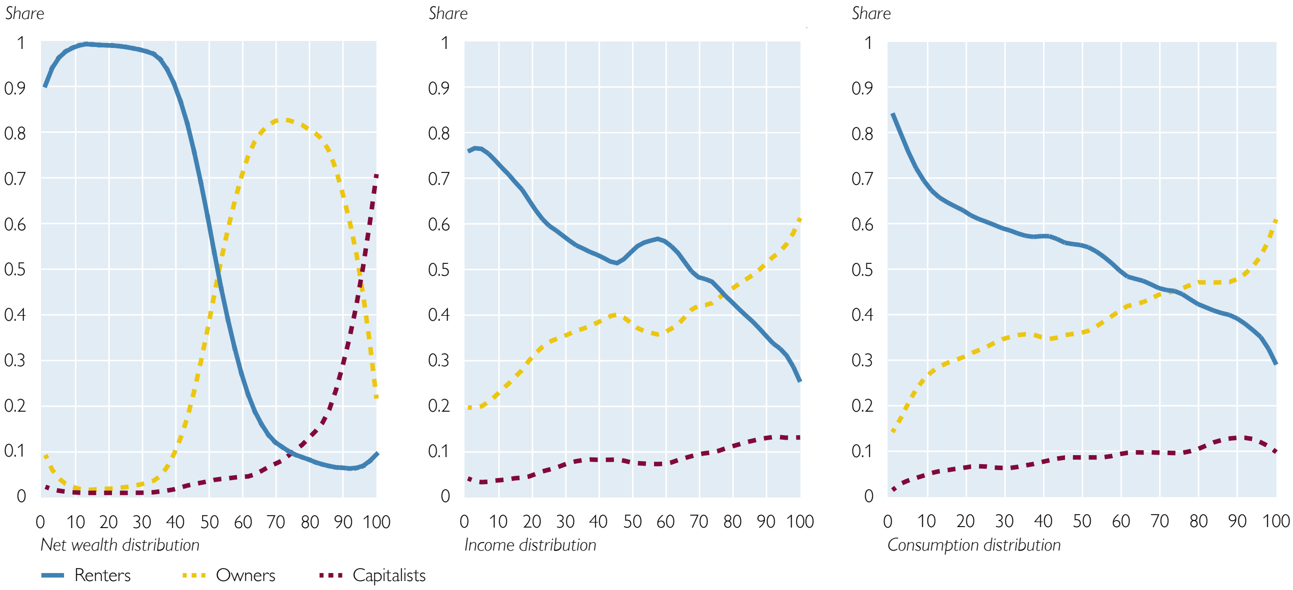

Source: HFCS 2017, OeNB; EU-SILC, HBS, Statistics Austria.

these three groups of income and consumption are remarkably similar. Only net wealth displays a different pattern.

For a more precise discussion of the multidimensional distribution of wealth, income and consumption for each of the groups of households, we calculate a similar chart as chart 1 for the three household groups. Due to the small number of observations on some parts of the distribution, we refrain from including the chart in the paper. We discover that renters are to be found over all parts of the combination of all 2D distributions. In the consumption-income space, owners and capitalists also occupy most of the cells. Thus, in each combination, there are households that have high, middle and low income connected to different consumption levels. For owners and capitalists, however, the space using also net wealth shows that at the bottom of the net wealth distribution, there are few owner households and almost no capitalists.

Table 4

Basic indicators for household debt burden

	Renter	Owner	Capitalist
	%		
Share of households holding debt (HFCS)	25	38	55
	Ratio		
Debt-to-asset ratio (HFCS)	7.7	0.3	0.1
Debt-to-(gross) income ratio (HFCS, EU-SILC)	0.3	2.6	1.6
	%		
Share of debtors with negative financial margin (HFCS, EU-SILC, HBS)	24	23	23

Source: HFCS 2017, OeNB; EU-SILC, HBS, Statistics Austria.

4.3 Household debt burden over three groups of households

Table 4 provides an overview of indicators generally used in the analysis of household debt burden and based on the matched data for each group of households.

The share of households holding debt is highest among capitalists, but the debt-to-asset ratio is lowest in this group and highest for renters. The debt-to-income ratio is highest for owners, which indicates a higher level of debt. Financial fragility – measured by the share of debtors with a negative financial margin – is roughly the same for all groups of households conditional on having debt. A deeper discussion of this topic going toward the estimation of potential exposure and loss given default is left for future research.

4.4 Robustness of the tails of the joint distribution

The tails of the distribution are of particular interest for many topics addressed with micro data. For the analysis of household vulnerability, we consider the left tail of the distribution. In the literature, analyses often also focus on the top tail of the distribution. This is why we concentrate on these parts and analyze the impact of the statistical matching procedure on them.

To this end, we look at the top and bottom 10% according to each indicator.[12] Extending the distribution by one or two dimensions, we analyze how many households remain in the tail of the joint distribution. Essentially, we investigate the cross-combination of belonging to the top 10% in each of the underlying one-dimensional distributions and show the share of households remaining in the top 10%. If all households in the top 10% of the net wealth distribution are also in the top 10% of the income distribution and the top 10% of the consumption distribution, we would see that 100% remain in the 2D and 3D distribution. We complement this analysis of results from the HFCS with the results from the matched data as well as another less granularly stratified single random hotdeck matching procedure. For the “matching II” procedure, we exclude ownership status of the household main residence and square meter of real estate and reduce age (EU-SILC as with

Chart 3

Top of the joint distribution and matching data sources

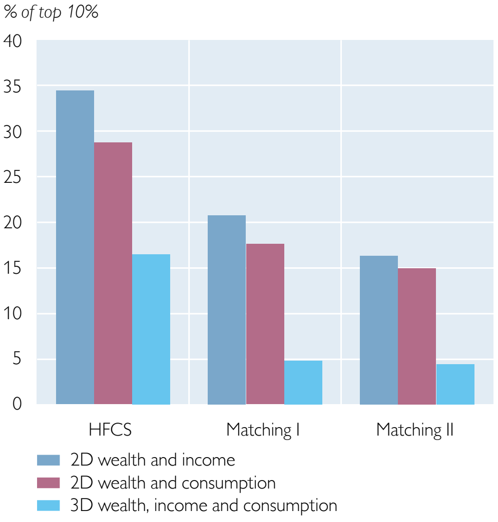

Source: HFCS 2017, OeNB; EU-SILC, HBS, Statistics Austria.

Chart 4

Bottom of the joint distribution and matching data sources

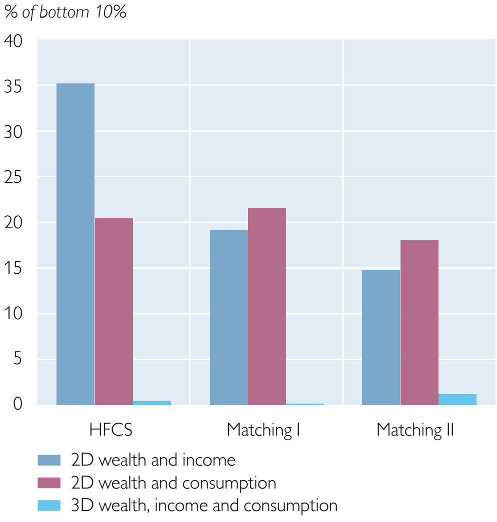

Source: HFCS 2017, OeNB; EU-SILC, HBS, Statistics Austria.

[12] *We tried to look at the top and bottom 5%, however, there are almost no households left when we look at the joint distribution. Thus, we turned toward the top and bottom 10%.*

HFCS) and household type to 3 categories (single, couple without children, rest) in the stratification.

Chart 3 splits the results in three blocks. First, we look at the information based solely on the HFCS, second, we use data from matching I, and third, we use a reduced-form matching procedure II. The first combination of wealth and income shows that from the 10% in each of the distributions of wealth and income, there are about 35% of households left (dark blue bar). The connection of wealth and consumption has under 30% of the top 10% in each category left. Taking the 3D distribution of wealth, income and consumption (light blue bar) together, only about 15% remain at the top in each indicator. Thus, increasing the dimensions of inequality reduces the share of households belonging to the top.

Taking the matched data in the two other blocks, we find an even larger reduction of the share of households remaining at the top of the multidimensional distribution. Therefore, depending on the specifics of the matching process, the procedure itself reduces the tail of the distribution. These results are deemed highly problematic when the data are used to discuss issues of inequality. Any statistics referring to matched data must be used with great caution.

Secondly, we turn toward the left, i.e. the bottom, tail of the distribution (see chart 4). The bottom of the distribution displays a similar picture as the top tail. The 3D distribution leaves almost nobody at the bottom of the distribution in each single dimension. Even if households have no wealth, they might have enough income to be lifted out of the bottom 10%. At the bottom tail of the 3D distribution, there are even fewer households left than at the top. The results show that we have to be very careful when analyzing household vulnerability and financial fragility. This is the reason why several indicators for vulnerable households are used.

The main reason for the impact of matching on the results seems to be the loss in correlation between the variables due to the random nature of matching. In order to investigate this result further, we show simple pairwise correlations of each of the involved indicators in the analyses in table 5. The correlation between income and wealth in the HFCS is about 0.7 and thus more than 6 times higher than the correlation of wealth from the HFCS and the matched income information from the EU-SILC, which is 0.1. Though less drastic for the correlation between wealth and consumption, there is still a factor of almost 2 comparing the matched information with variables contained in one dataset. A similar difference is found between the income and

Table 5

Correlation matrix

	HFCS			EU-SILC	HBS
	Net wealth	Gross income	Consumption	Gross income	Consumption
	EUR thousand				
HFCS					
Net wealth	1.00				
Gross income	0.65	1.00			
Consumption	0.24	0.35	1.00		
EU-SILC					
Gross income	0.10	0.09	0.15	1.00	
HBS					
Consumption	0.14	0.12	0.21	0.20	1.00

Source: HFCS 2017, OeNB; EU-SILC, HBS, Statistics Austria.

consumption correlation within the HFCS (0.4) compared to the correlation between income and consumption using both matched information (0.2).

5 Conclusion

The conceptual challenges of a multidimensional and relational approach in understanding inequality are far-reaching and call for caution.

In our empirical results, we find a tendency of the statistical matching process toward the mean. Additionally, the issue of estimating uncertainty is largely unresolved in the literature. A closer inspection of this and differences across countries, however, is left for future research.

The distributional information covered in the HFCS is limited but it is unbiased and should be sufficient to base scientific discussions of the joint distribution on this data base. The correlation between wealth, income and consumption is substantially higher in the HFCS data than in the statistically matched data. Research on financial stability is already taking into account the multidimensional nature of inequality but the unresolved conceptual challenges remain manifold.

References

Albacete, N., S. T. Dippenaar, P. Lindner and K. Wagner. 2018. Eurosystem Household Finance and Consumption Survey 2017. Methodological notes for Austria. In: Monetary Policy & the Economy Q4/18 – Addendum. Vienna: OeNB.

Albacete, N., P. Lindner and K. Wagner. 2016. Eurosystem Household Finance and Consumption Survey 2014. Methodological notes for Austria. In: Monetary Policy & the Economy Q2/16 – Addendum. Vienna: OeNB.

Albacete, N., P. Lindner, K. Wagner and S. Zottel. 2012. Eurosystem Household Finance and Consumption Survey 2010. Methodological notes for Austria. In: Monetary Policy & the Economy Q3/12 – Addendum. Vienna: OeNB.

Albacete, N. and P. Lindner. 2013. Household Vulnerability in Austria – A Microeconomic Analysis Based on the Household Finance and Consumption Survey. In: Financial Stability Report 25. Vienna: OeNB. 57–73.

D'Orazio, M., M. Di Zio and M. Scanu. 2006. Statistical Matching. Theory and Practice. New York: Wiley.

Eurostat. 2013. Statistical Matching: a model-based approach for data integration. Methodologies and Working papers. Luxembourg: Eurostat.

Fessler, P., P. Lindner and M. Schürz. 2018. Eurosystem Household Finance and Consumption Survey 2017 for Austria. In: Monetary Policy & the Economy Q4/18. Vienna: OeNB. 36–66.

Fessler, P., P. Lindner and M. Schürz. 2018. Eurosystem Household Finance and Consumption Survey 2017 – first results for Austria. In: Monetary Policy & the Economy Q4/18. Vienna: OeNB. 36–66.

Fessler, P., P. Lindner and M. Schürz. 2016. Eurosystem Household Finance and Consumption Survey 2014 – first results for Austria (second wave). In: Monetary Policy & the Economy Q2/16. Vienna: OeNB. 34–95.

Fessler, P., P. Mooslechner and M. Schürz. 2012. Household Finance and Consumption Survey of the Eurosystem 2010 – first results for Austria. In: Monetary Policy & the Economy Q3/12. Vienna: OeNB. 24–62.

Fessler, P. and M. Schürz. 2018. The functions of wealth: renters, owners and capitalists across Europe and the United States. OeNB Working Paper 223.

Fisher, J., D. Johnson, T. Smeeding and J. Thompson. 2018. Inequality in 3-D: Income, Consumption, and Wealth. Finance and Economics Discussion Series 2018-001. Washington: Board of Governors of the Federal Reserve System.

Lamarche, P. 2017. Measuring Income, Consumption, and Wealth jointly at the micro-level. Luxembourg: Eurostat.

Leulescu, A. and M. Agafitei. 2013. Statistical matching: a model based approach for data integration. Eurostat Methodologies and Working Papers.

Meyer, B. D., W. K. C. Mok and J. X. Sullivan. 2015. Household Surveys in Crisis in Journal of Economic Perspectives – Volume 29. Number 4. Fall 2015. 199–226.

OECD. 2013. OECD Framework for Statistics on the Distribution of Household Income, Consumption and Wealth. Paris: OECD Publishing.

Piketty, T. 2014. Capital in the Twenty-First Century. Cambridge, Massachusetts: The Belknap Press of Harvard University Press.

Pistor, K. 2019. The Code of Capital. How Law creates wealth and inequality. Princeton University Press.

Statistics Austria. 2018a. Standard-documentation – Meta information (Definitions, comments, methods, quality) on EU-SILC 2017.

Statistics Austria. 2018b. Standard-documentation – Meta information (Definitions, comments, methods, quality) on Household Budget Survey 2014/15.

Statistics Austria. 2016. Standard-Dokumentation Metainformationen (Definitionen, Erläuterungen, Methoden, Qualität) zu EU-SILC 2014.

Statistics Austria. 2013. Standard-documentation – Meta information (Definitions, comments, methods, quality) on Household Budget Survey2009/10.

Statistics Austria. 2012. Standard-Dokumentation Metainformationen (Definitionen, Erläuterungen, Methoden, Qualität) zu EU-SILC 2010.

Stiglitz, J. E., A. Sen and J. P. Fitoussi. 2009. Report of the Commission on the Measurement of Economic Performance and Social Progress. https://ec.europa.eu/eurostat/documents/118025/118123/Fitoussi+Commission+report.

Annex

Technical aspects of the matching process

Statistical matching does not come in a single agreed upon method. There are the following approaches that are commonly taken when statistically matching disjoint data sets that cannot be linked with a common identifier:

- Regression-based parametric approach
- (stratified) random hotdeck
- rank hotdeck
- reweighting

The first method involves a regression of the desired information that is missing in the dataset of choice – say A – but is included in another dataset, e.g. income in our case, on a set of common variables in both datasets. The coefficients of the regression are transferred and applied to the values in dataset A, thereby forecasting the missing information. In the hotdeck procedures, a random donor household is chosen and its value attached to the receiver data. This random process is complemented by a strategy to find similar households, i.e. stratification or ranking along a certain distribution. In the stratification process, a set of information is used to create bins certain households belong to because they fulfill characteristics from the chosen parameters. Depending on the information, the stratification can be more or less granular, meaning there are many bins with few similar households or there are few bins with many households. If there are no households from the donor data set in a bin, the households cannot be assigned a value in the receiver dataset in this bin. This leads to bias, since these households ultimately have to be left out of the analysis. In the "rank hotdeck" procedure, the characteristics defining the bin are whether they belong to a certain rank in a chosen distributional dimension. In the univariate case with only one distribution, there are no missing households from the donor dataset, since the distribution in each dataset is complete. The choice of the distribution, however, is crucial and a single dimension might not be enough to guarantee conditional independence. In the last method, the weights of households are adapted in such a way that the distribution of the desired outcome variable, like income, follows what is observed in the donor data. This approach requires some similar information in the receiver dataset. As explained in the main text, we opted for a stratified single random (rank) hotdeck procedure.

In practice, the stratified single random hotdeck procedure requires at least the following steps to generate the desired dataset:

1. select donor and receiver datasets
2. select stratification variables
3. select variables from the donor data that should be matched to the receiver data
4. stratify your sample
5. randomly select a donor household for a receiver household

Each of the above steps requires assumptions that potentially affect the results and the variance estimation around a specific result. Far from giving a complete in-depth analysis what each and every assumption entails, we want to lay out here which assumptions we have made in order to provide for transparency in the choices made and make it possible to discuss alternatives.

First, as introduced in the main text in table 1, we base our analysis on the three main surveys in Austria. The HFCS covers information on wealth, income and consumption. The EU-SILC focuses on income and the HBS on consumption. Although the HFCS is arguably the most complete survey, details on consumptions are covered to a larger extent in the HBS, and the EU-SILC has access to register income data, which should improve the quality of the data provided. We consider both EU-SILC and HBS data as the donor datasets for income and consumption, respectively. This means our receiver dataset is the HFCS. This choice is largely due to the issue that it is even harder to match wealth information to another dataset compared to matching any other information to the data containing wealth. However, apart from this, the decision is ad hoc and results would change if it is altered. In our case the decision is made more complicated by the time dimension. The most obvious choice is to use the same data sources as receiver and donor datasets over time. That is what we have done. However, considering some information on the quality of the matching process, e.g. how many households ultimately end up remaining missing, alterations to the receiver and donor datasets could enhance quality. Additionally, since there are three waves of the HFCS, we looked for information from the EU-SILC and HBS in the respective time period. While this seems relatively unproblematic for matching with EU-SILC data, there are serious complications in the matching process with data from the HBS, which is conducted only every 5 years. The first consumption information was provided one to two years earlier than the information provided in the HFCS, and the 2014 wave of the HBS has to be used twice for both the second and the third waves. Although this implies limitations in the results, there is nothing that can be done about it apart from urging the data provider to provide for a common survey period.

Table A1

Stratifiction variables

		Matching I			Matching II		
Matching information	Description	Number of categories	Coding	Used in the matching of	Number of categories	Coding	Used in the matching of
Age	Age of household head	8	Up to 19 20 to 29 30 to 39 40 to 49 50 to 59 60 to 69 70 to 79 80 and above	HFCS-EU-SILC	× × × × × × × ×	× × × × × × × ×	× × × × × × × ×
	Age of household head	3	Up to 29 30 to 59 60 and above	HFCS-HBS	3	Up to 29 30 to 59 60 and above	HFCS-EU-SILC HFCS-HBS
Education	Level of education of household head	3	Without secondary education Secondary education Tertiary education	HFCS-EU-SILC HFCS-HBS	3	Without secondary education Secondary education Tertiary education	HFCS-EU-SILC HFCS-HBS
Household structure	Type of household classified with respect to age of household members and their relationship	6	Single – no children Couple – no children More than 3 adults – no children Single with children Couple with children Three or more adults with children	HFCS-EU-SILC HFCS-HBS	3	Single Couple – no children Other	HFCS-EU-SILC HFCS-HBS
Tenure status	Ownership structure of main resi-dence	2	Owner (including free usage) Renter	HFCS-EU-SILC HFCS-HBS	× ×	× ×	× ×
Income	Disposable household income	5	1st quintile 2nd quintile 3rd quintile 4th quintile 5th quintile	HFCS-HBS	5	1st quintile 2nd quintile 3rd quintile 4th quintile 5th quintile	HFCS-HBS
Wealth indicator	Size of main residence	5	1st quintile 2nd quintile 3rd quintile 4th quintile 5th quintile	HFCS-EU-SILC HFCS-HBS	× × × × ×	× × × × ×	× × × × ×

Source: Authors' compilation.

As regards the choice of stratification variables, for most of the results in the paper we opt for a relatively fine stratification. This implies that matched households are relatively similar along the dimensions age, education, household structure, tenure status, income (only for HFCS-HBS matching) and the size (in square meters) of the households' main residence. Table A1 shows the details. The categories of the household structure follow a logical separation in households with and without children followed by households with 1, 2 and 3+ persons. Underlying this information, there are additional assumptions, we have to make along the way. As a start, the details of the definition of a household and the coverage of the total household population is not identical over datasets. The HFCS, e.g., also collects information on households that are not in the register of residents. Another example, the way how the data are collected, and the definition of income are not completely identical in the HBS and HFCS. Also, net income is available only starting from wave 2 in the HFCS, so in wave 1 gross income is used instead. Another example is that floor space is top coded in the EU-SILC but not in the HFCS, which limits what can be done in terms of percentiles. It would be preferable to ensure harmonized definitions over the underlying datasets. This should be at the core of further improving the surveys. There is also an implicit assumption of whether to hold stratification stable over time or change it appropriately in each wave. We opted for a constant stratification over time in order to ensure that changes over time are not affected by this choice. There is an indefinite number of imaginable breakdowns for stratification. In order to be able to discuss the impact of changes in the assumption about stratification, we discuss in the paper an exercise where we simplify our matching procedure to what we call "matching II." As shown in table A1, this essentially reduces the number of similar characteristics households need to have for being matched. Thus, it reduces the number of observations that cannot be matched and hence remain missing, but this comes at the cost that households that are matched now are more dissimilar. Thus, there is a trade-off between the two ways of stratification and hence a priori no clear advantage of one over the other. The distribution of households in these stratification categories should be similar over the underlying datasets to ensure a smoothly working matching procedure. This similarity is guaranteed only for quintiles in which there are, by definition, 20% of the household population. Some of the other variables are used in the post-stratification process to reach final household weights and each, in turn, is based on the micro census and hence should be relatively similar. Due to space constraints, we refrain from reporting tests for the similarity of the distribution, such as Hellinger distance, but instead perform sensitivity analyses with respect to different matchings.[13]

Table A2 gives an overview of the practical implications of the stratification introduced above. It shows the number of distinct bins in the matching process between all datasets. Additionally, we see how many bins are occupied to what extent and what the missing rate pattern in the matched data is. The first two columns headed "matching I" provide the information for the finely stratified procedure that underlies most of the results in the paper.

There are 1,440 and 2,700 possible distinct bins in matching HFCS with EU-SILC and HBS data, respectively. This number does not change over time, since the

[13] *For more details on the testing of similarity of the underlying distributions of households see e.g. Leulescu and Agafitei (2013).*

Table A2

Number of matching strata

	Matching I		Matching II	
	Matching HFCS EU-SILC	Matching HFCS HBS	Matching HFCS EU-SILC	Matching HFCS HBS
Number of possible strata (bins)	1,440	2,700	27	135
Wave 2017				
Occupied strata in the HFCS	591	819	27	119
Occupied strata in donor data	725	1.039	27	129
Number of strata in HFCS without any complete cases in donor data	69	163	0	2
Number of strata with only 1 complete case	183	348	0	3
Number of strata with only 2–5 complete cases	241	356	0	19
Approximate number of households with missing matched information	83	162	0	2
Wave 2014				
Occupied strata in the HFCS	589	843	×	×
Occupied strata in donor data	754	1.039	×	×
Number of strata in HFCS without any complete cases in donor data	65	182	×	×
Number of strata with only 1 complete case	202	348	×	×
Number of strata with only 2–5 complete cases	261	356	×	×
Approximate number of households with missing matched information	73	169	×	×
Wave 2010				
Occupied strata in the HFCS	589	937	×	×
Occupied strata in donor data	744	1.100	×	×
Number of strata in HFCS without any complete cases in donor data	79	208	×	×
Number of strata with only 1 complete case	183	362	×	×
Number of strata with only 2–5 complete cases	273	423	×	×
Approximate number of households with missing matched information	87	160	×	×

Source: HFCS Austria 2017, OeNB, Statistics Austria, EU-SILC 2016 and HBS 2014/15.

Note: The number of households with missing matched information is the average over the five implicates and thus can only be given as an approximation.

categories used in the stratification are not altered either. The number of occupied cells, however, does change over time and is about 590 and from 820 to 940 for the HFCS in the matching with EU-SILC and HBS data, respectively. In the EU-SILC, the number of occupied cells ranges from 725 to 754 and for the HBS from 1,039 to 1,100. Due to the sample size, the number of occupied cells in the donor datasets is larger than in the receiver dataset. Although this is desirable, it is not sufficient to ensure that all households in the HFCS actually have potential matches. Depending on the indicator and the wave, there are between 70 and 170 households that cannot be matched and remain missing in the desired information. This obviously introduces bias in the estimation. There is thus a trade-off between the precision of the stratification and the resulting number of missing observations. Inspecting the last two columns headed "matching II," where we simplified the stratification, we see almost no missing observations reported. To repeat, the cost of that is that the matched households are not very similar in many dimensions and thus the CIA assumption is hardly met.

Table A3

Shares held by groups in society

	Population share	Net wealth	Gross income	Consumption
	%			
Renter	54	32	47	47
Owner	38	54	43	43
Capitalist	8	14	10	10

Source: HFCS 2017, OeNB; EU-SILC, HBS, Statistics Austria.

Furthermore, table A2 shows the number of cells where there are very few households in the donor dataset. This number ranges from about 180 to 360, depending on the survey wave and donor for cases, where there is only one household to be matched to potentially many HFCS households. The randomness of the matching procedure reduces to this single household being selected with certainty. Even in the case where two to five households are in a specific bin in the donor dataset (240 to 420 cases), the likelihood of one household being selected is high. Comparing this information to "matching II" shows the trade-off between the similarity of households to be matched and the randomness in the process. Any choice has implications for the results and, to the best of our knowledge, there is no agreement in the literature on the optimal choices to be made.

For us it is relatively clear which information to be matched from which database, and thus, steps three to five are straightforward to implement. We use the user-written Stata command "hotdeck" for the matching procedure.[14] As the data of the HFCS are multiply imputed (see e.g. Albacete et al., 2018), one complication arises in the choice of how to match households across implicates. This implicitly means that we have to assume either that one household is matched with the same household from the donor dataset over all implicates or that each household in each implicate is treated as separate unit and thus can be matched with different households across implicates. We opted for the latter since some information from the stratification might be imputed and thus change over implicates and hence falls in a different bin in the matching process. There is also the possibility to use multiple matches for one household, similar to multiple imputation. This, however, would extensively enlarge the dataset and thus is disregarded in the analysis at hand. Nonetheless, all the assumptions might have an impact on the results.

Shares of groups of households in society

Table A3 shows the share of each group of households in society. More than half of the population rents their main residence, only 8% are capitalists. The remaining 38% are classified as owners. We see that in terms of wealth, renters hold a disproportionately small share while owners and capitalists hold a disproportionately large share. These results are far less pronounced – though still visible – for income and consumption.

Additional information at the top and bottom of the distribution

While the main results are portrayed in the paper, table A4 looks at the top of the distribution and how shares are impacted by the extension to more than one dimension and by the matching procedure. For each of the possible three 2D distributions and the 3D distribution, we show the share of households remaining at the top of the distribution.

Table A5 provides similar information for the left tail of the distributions.

[14] *The command was written by Adrian Mander, MRC Human Nutrition Research, Cambridge, U.K.*

Table A4

Impact of matching on the share of households at the top of the distribution

		HFCS			Matching I			Matching II		
		Net wealth	Income	Con-sumption	Net wealth	Income	Con-sumption	Net wealth	Income	Con-sumption
		%								
Single indicator	**Share of households**	10.0	10.0	10.0	10.0	10.0	10.0	10.0	10.0	10.0
2-dimensional distribution	Net wealth and income	**3.4**			**2.1**			**1.6**		
	One without the other	6.5	6.5		7.9	7.6		8.3	8.3	
	Net wealth and consumption	2.9		2.9	1.8		1.8	1.5		1.5
	One without the other	7.1		7.1	8.2		7.7	8.5		8.5
	Income and consumption		**3.7**			**1.9**			**1.7**	
	One without the other		6.3	6.3		7.9	7.6		8.3	8.3
3-dimensional distribution	Share of households	**1.6**			**0.5**			**0.4**		
	Net wealth and income but not consumption	**1.8**			**1.6**			**1.2**		
	Net wealth and consumption but not income	1.2		1.2	1.3		1.3	1.1		1.1
	Income and consumption but not net wealth		**2.1**			**1.4**			**1.2**	

Source: HFCS 2017, OeNB; EU-SILC, HBS, Statistics Austria.

Table A5

Impact of matching on the share of households at the bottom of the distribution

		HFCS			Matching I			Matching II		
		Net wealth	Income	Con-sumption	Net wealth	Income	Con-sumption	Net wealth	Income	Con-sumption
		%								
Single indicator	**Share of households**	10.0	10.0	10.0	10.0	10.0	10.0	10.0	10.0	10.0
2-dimensional distribution	Net wealth and income	**3.5**			**1.9**			**1.5**		
	One without the other	**6.5**			8.1	7.9		8.6	8.6	
	Net wealth and consumption	2.0		2.0	2.2		2.2	1.8		1.8
	One without the other	8.0		8.0	7.9		7.4	8.2		8.2
	Income and consumption		**2.3**			1.6	1.6		**1.3**	
	One without the other		7.7	7.7		8.2	7.9		8.8	8.7
3-dimensional distribution	Share of households	**0.0**			**0.0**			**0.1**		
	Net wealth and income but not consumption	**0.2**			**0.3**			**0.6**		
	Net wealth and consumption but not income	0.5		0.5	0.5		0.5	0.6		0.6
	Income and consumption but not net wealth		**2.3**			1.6	1.6		**1.2**	

Source: HFCS 2017, OeNB; EU-SILC, HBS, Statistics Austria.