Eurosystem Household Finance and Consumption Survey 2023: Methodological notes for Austria

The Eurosystem Household Finance and Consumption Survey (HFCS) collects data on households' wealth, income and consumption across the euro area. This report describes how households were selected and interviewed, and how the data were processed in the 2023 survey wave in Austria. Its main objective is to provide transparent documentation of all the steps in data production, such as careful editing and weighting, and the statistical procedures used to ensure high data quality.

Authors

Nicolas Albacete
OeNB, Research Section
nicolas.albacete@oenb.at

Peter Lindner
OeNB, Research Section
peter.lindner@oenb.at

Representative household microdata for Austria

The report documents in detail how around 3,000 households were selected and weighted to ensure that the survey produced reliable and meaningful results representative of all Austrian households.

High data quality through strict standards

Missing information was supplemented by using advanced techniques (multiple imputation); more than 250 consistency checks ensured valid responses. All changes to the data made during these procedures are well documented and fully traceable.

Foundation for research and policymaking

This methodologically sound survey provides key data that feed into analyses of monetary policy, financial stability and consumption, both nationally and across euro area countries.

Contents

The HFCS in Austria (fifth wave): key facts

Methodological framework at a glance

Questionnaire

The HFCS in Austria is based on an internationally harmonized questionnaire that covers the key stock and flow components of the household balance sheet and integrates them with socioeconomic characteristics. Data were collected from households.

Reference period

The data on stock positions and socioeconomic characteristics refer to the survey date (the fieldwork was carried out between late March 2023 and February 2024). Income-related data refer to the 2022 calendar year. The data on household consumption refer to a typical month.

Geographical scope

Austria

Sampling

Target population

All households in Austria (irrespective of nationality and citizenship)

Sampling frame

Postal addresses of all households in Austria

Sampling design

Stratified two-stage cluster sample design
- Stratification: NUTS-3 regions divided into 8 classes by municipality size
- Primary sampling unit (PSU):   enumeration districts
- Secondary sampling unit (SSU): postal addresses

The gross sample comprised a total of 689 PSUs and 7,636 SSUs in 192 strata.

Survey company

Institut für empirische Sozialforschung GmbH (IFES)

Fieldwork

General information

Fieldwork period: March 2023 to February 2024
Number of interviewers: 48
Method of data collection: computer-assisted personal interview (CAPI)

Interviewer training

Number of (virtual) HFCS interviewer training sessions: 7 training sessions, designed as full-day training over two days if held virtually

Pilot survey

Number of pilot interviews: 50

Contact strategy

All households received a personalized letter from the Governor of the Oesterreichische Nationalbank (OeNB) and an information leaflet distributed by IFES before they were contacted by the interviewers.

The interviewers had instructions to make up to five contact attempts per household over a period of at least three weeks. At least two of these attempts were to be made in person, at least one attempt was to be made on a weekend and another outside regular working hours (9:00 a.m. to 5:00 p.m.).

Incentives for participation

Participation in the HFCS was voluntary.

Each household that successfully completed an interview received a voucher with a face value of EUR 15.

Documents used during interviews

Showcards, interviewer manual, glossary

Interviewer monitoring

The survey company monitored the conduct of the interviews by randomly checking (by telephone) around 15% of the interviews.

The (anonymized) data from the completed household interviews were forwarded to the OeNB in 24 batches during the field phase to enable the prompt assessment of each interview and interviewer.

Follow-up queries by telephone

Follow-up queries were made by telephone to address outliers and inconsistent responses from some 200 households.

Editing and consistency checks

Number and type of edits

Number of observations: around 1.4 million
Thereof edited observations: around 44,600 (about one-third on the basis of verbatim records)
Percentage of edited observations: 3.1%

Consistency checks during the interviews

Number of consistency checks
programmed into the questionnaire: around 250

Postinterview consistency checks

Expert analysis of the data from each interview, follow-up phone calls to clarify uncertainties, investigation of outliers and consistency checks of the information collected, technical review of filtering.

Flag variables are used to document all edits and imputations.

Imputations

Method: multiple imputation with chained equations (broad conditioning approach)

Number of multiple imputation samples:	5
Number of iterations per imputation sample:	10
Median of the variables with missing values per household:	10
Mean of the variables with missing values per household:	16.7

Sample size and response rate

Number of households in the sample (gross sample):	7,636
Number of successfully interviewed households (net sample):	2,849
Number of households that could not be reached (despite five contact attempts):	75
Number of households that refused to participate:	4,462
Number of households that did not participate for other reasons:	164
Number of addresses whose eligibility was unknown:	1
Number of ineligible addresses:	85
Incomplete interviews and interviews discarded after fieldwork:	20
Response rate:	38%

Weighting

Final weights computed with nonresponse and poststratification adjustments to design weights – methods:

- Model-based adjustment combined with weighting-class adjustment, based on optimal number of classes (nonresponse adjustment)
- Cell adjustment (poststratification adjustment)

Smallest final weight:	319
Median of final weights:	1,268
Mean of final weights:	1,437
Highest final weight:	8,561
Sum of final weights (target population):	4,095,350
Unequal weighting effect:	1.261

No trimming and no normalization of weights

Variance estimation

Method:	Rescaling bootstrap procedure
Number of replicates:	1,000
Number of pseudo-strata:	151

Computation of replicate weights: adjustments made to design weights to obtain replicate weights are identical to adjustments made to obtain the final weights.

Finite population corrections were applied to all replicate weights.

1 Introduction

The Eurosystem Household Finance and Consumption Survey (HFCS) is the most comprehensive compilation of data capturing real assets, financial assets, debt and expenditures of households in one survey, allowing for in-depth scientific analyses of household balance sheets in line with international standards. HFCS data are comparable across all participating countries thanks to the ex ante harmonization of the survey and of the survey methods applied. The geographical scope of the HFCS has been broadened with every wave, to include more than 20 countries in the fifth wave. In Austria, the euro area HFCS was first carried out in 2010/11, a second time in 2014/15, a third time in 2016/17, a fourth time in 2021/22 and most recently in 2023/24. All five waves were conducted by the Oesterreichische Nationalbank (OeNB) in cooperation with the survey company IFES (Institut für empirische Sozialforschung GmbH). The ECB is expected to make data for all participating countries from the fifth HFCS wave available for research purposes in summer 2026.

This publication provides an in-depth view of the data collection process and the methods applied. Based on the methodological documentation from the first four HFCS waves in Austria (Albacete et al., 2012, 2016, 2019, 2023), it aims at making the process of data collection as transparent as possible and serves as the basis for correct evaluation of HFCS data. Since the first wave, specific methodological aspects have been discussed in a number of publications. For instance, the information gathered from respondents' verbatim answers (Lindner and Schürz, 2017) and from the survey of interviewers has been examined in depth and cross-checked with the HFCS data (Albacete et al., 2025; Albacete and Schürz, 2013b and 2015). Other papers have discussed the relevance of paradata and ways to improve them (Albacete and Schürz, 2014a and 2014b) as well as comparability with other surveys conducted in Austria (Albacete and Schürz, 2013a) and across HFCS countries (Andreasch et al., 2013). Moreover, different approaches to compiling the components of the household balance sheet have been compared (Lindner and Schürz, 2015) and methodological enhancements between the first and the second wave of the HFCS in Austria have been discussed (Lindner et al., 2014).

The chapters are self-contained, each dealing with specific aspects of the HFCS, and can therefore be read independently of each other. Cross-references help the reader recognize links to other chapters or material aspects discussed within them. The sequence of chapters reflects the logical flow of the survey. Closely related topics (e.g. constructing survey weights and estimating the correct variance using HFCS data) are arranged in a way to ensure comprehensibility. To avoid redundancies, only essential details were repeated. The following seven chapters provide a detailed explanation of each step in the survey process, with a further chapter designed as a user guide.

Chapter 2 on the questionnaire of the HFCS in Austria explains the content of the survey, discussing the individual parts and special features of the questionnaire, the sequence of questions as well as the unit of data collection.

Chapter 3 looks into the role of the interviewers who conducted the face-to-face interviews. Great importance was placed on the qualifications of the interviewers as well as on their professional demeanor and expertise, which all contributes significantly to the quality of data obtained. The chapter also covers details on the contact strategy and incentives for households to participate in the survey. Moreover, it outlines the information material and documents that were made available to the households in the HFCS sample.

All raw data collected by the interviewers were reviewed during the field phase, leading to the collection of further data or data edits where necessary. This process is described in detail in chapter 4 on consistency checks and editing, which lists all changes to the raw data as well as the flags included in the dataset to document such changes.

Chapter 5 on the multiple imputations applied in the HFCS deals with item nonresponse. For cases in which respondents were unwilling or unable to answer one or several questions, we performed multiple imputations to obtain the missing information. This approach made it possible to correct distortions due to item nonresponse at least to some extent and also to account for the uncertainty attached to imputations, which, like all edits, have been flagged. Users of the HFCS data may apply our imputations or deal with item nonresponse in a different way.

Chapter 6 on sampling provides a detailed description of the survey sample design. The complex survey sample design used for the first wave was developed further for the second wave to ensure a sufficiently representative sample of Austrian households that fits the purpose of the Eurosystem and the OeNB. From the third wave onward, we have generally retained the sample design methodology of the second wave.

The final household weights were calculated in several steps on the basis of the sampling design. Chapter 7 outlines the procedure for the construction of survey weights. The sampling design yields design weights for each household already in the sampling process. It takes several steps to process these weights to account for information obtained during the field phase (such as nonparticipation of households and external information regarding the distribution of certain household characteristics).

Another step is required to obtain the correct variance estimation, which is presented in chapter 8 on the construction of replicate weights for variance estimation.

The user guide in chapter 9 provides basic guidance on the correct use of HFCS data in Stata.

Finally, an online appendix (https://www.hfcs.at/en/results-tables/) serves to provide all the essential documentation and background material used in the HFCS (available in German only). The HFCS website also provides information about the publication of HFCS data from all participating countries by the ECB (expected in summer 2026) and any other HFCS news.

2 Questionnaire

2.1 Introduction

Almost twenty years ago, in 2006, the Governing Council of the European Central Bank (ECB) decided to establish the Household Finance and Consumption Network (HFCN) to collect data on wealth, income and consumption from a representative sample of households. The ultimate goal of the network is to ensure ex ante harmonization of the information collected. Harmonization was also broadly achieved with regard to methodology. Nevertheless, some cross-country differences in the technical conditions for the HFCS surveys, for example the respective sampling designs (see chapter 6) or multiple imputation procedures (see chapter 5) remain.

It is the intention to conduct the HFCS every three years. In Austria, the survey has now been conducted five times, including the fifth wave documented here (HFCS 2010, HFCS 2014, HFCS 2017, HFCS 2021 and HFCS 2023). The questionnaire for the fifth HFCS wave was based on the questionnaires used and experiences made in the first four waves. This chapter presents the Austrian questionnaire, which was designed on the basis of the HFCN blueprint questionnaire (drafted in English)[1], expanded to include country-specific features (e.g. loan specifics or the national variant of housing association apartments).

This chapter is structured as follows: First we outline the objectives of the HFCS survey in Austria (section 2.2), define the data collection unit (section 2.3), and the reference period (section 2.4). Then we describe the sequence of the questionnaire and highlight some core questions and variables (section 2.5). We subsequently discuss special features of the questionnaire (section 2.6) and list interviewer documents (section 2.7) and participating countries (section 2.8). Section 2.9 refers to further information provided in an online appendix.

2.2 Objectives of the survey

The main goal of the HFCS is to collect microdata on the structure of the assets and liabilities of households in euro area countries as well as some EU countries that have not yet adopted the euro,[2] as these data allow for the analysis of households' investment and consumption decisions. These survey data can be used, for example, to:
- gain insights into various aspects of the monetary transmission mechanism and of financial stability,
- gain insights into individual household behavior,
- analyze the impact of economic policy measures and macroeconomic shocks, and to
- make cross-country comparisons.

Household-level finance and expenditure data are indispensable for a central bank, as they contribute significantly to improving the analysis of monetary policy and financial stability. Furthermore, the economic developments of the past decade and more have clearly shown that it is not the level of household debt – which can be calculated from macrodata – that matters most in the assessment of stability risks, but the specific burden on different income, occupational and age groups. Hence, decision-making on monetary policy and financial stability issues is also informed by analyses based on the HFCS, which is the

---

[1] For further details about the HFCN network and the HFCS survey (including the euro area blueprint questionnaire), see www.ecb.europa.eu/pub/economic-research/research-networks/html/researcher_hfcn.en.html (accessed on July 17, 2025).

[2] Some euro area countries (such as Ireland and Estonia) did not participate in the first wave (HFCS 2010). In the fourth wave, some EU countries (Hungary and Czechia) that have not adopted the euro also participated in the survey.

most comprehensive household-level survey conducted on this subject in the euro area. Publications attesting to the variety of insights gained from analyses relying on the microdata from the first four HFCS waves are, amongst others, Albacete and Lindner (2013, 2015, 2017a and 2017b), Albacete et al. (2014), Albacete et al. (2016 and 2018b), Albacete et al. (2018a), Beer and Wagner (2017) Fessler and Schürz (2013, 2017, 2022), Fessler et al. (2014, 2015 and 2017), Lindner and Redak (2017) and Wagner (2014), Lindner et al. (2022), Fessler and Schürz (2022), Drescher et al. (2020), Bekthiar et al. (2019), Albacete et al. (2025), Kennickell et al. (2021), Lindner and Schürz (2019, 2021), Lindner (2021), Albacete et al. (2020, 2021), Lindner and Probst (2020).

2.3 Unit of collection

To begin with, a survey requires a clear definition of the target population (see also chapter 6) and the data analysis unit. In the case of the HFCS, households represent the main unit of collection, although some data are also collected from individuals. The household definition is contained in the HFCS common template.

2.3.1 Definition of "household"

For the purpose of the HFCS, a household is defined as an individual living alone or a group of people living together in the same private dwelling who share household expenses and jointly take expenditure decisions.[3] More specifically, household members include:
- people who live in the same household and are related to each other,
- people who share household expenses and live in the same household but are not related,
- people who usually live in the same household (the reference period being the six months before the interview) but are temporarily absent e.g. because of holiday travel, job assignments away from home, hospital stays or boarding school stays, and
- children who are educated away from home but do not constitute a separate household, i.e. remain financially dependent on their family.

The household definition also includes people who have been members of the household for less than six months at the time of the interview (e.g. a new partner or a new child), provided that they share household expenses with the other members or are fully financially dependent on the household (children).

Employees of residents, like au-pairs or nursing staff, short-time visitors or subtenants are considered separate households. In shared apartments, all residents are treated as separate households unless they also share household expenses. This means that a particular address may be used by more than one household as defined for HFCS purposes (e.g. people sharing a residence). In such cases, we selected the household whose member received the letter of invitation to participate in the survey.

The definition also includes households with non-German speaking household members, households living at a place registered as their second home in the centralized residence registry and households not officially registered at a particular residence but living there.

---

[3] The HFCS household definition is discussed in greater detail in chapter 6. However, this section outlines the definition as is required for this chapter to be able to stand alone.

2.3.2 Financially knowledgeable person

All questions pertaining to the household were put to the person that the members of the household deemed to be most familiar with the household's finances, i.e. liabilities, assets, income and expenditure. This person, referred to as the financially knowledgeable person (FKP), answered all questions relating to the household as a whole (sand-colored sections in figure 1) as well as individual questions on behalf of absent household members. As a rule, questions relating to individuals were meant to be answered individually by all household members aged 16 and over.

Since the FKP was typically a member of the household, they were also the reference person of the given household. In line with the approach taken in the previous waves, the FKP could also be a third person who was not a household member. This could be a family member (e.g. a child) who oversaw the household's finances but was no longer a member of the household or a household's tax consultant or financial advisor. Only four households in the sample from the fourth wave had an FKP who was not a household member. In this case a member of the household was selected as the reference person for the household.

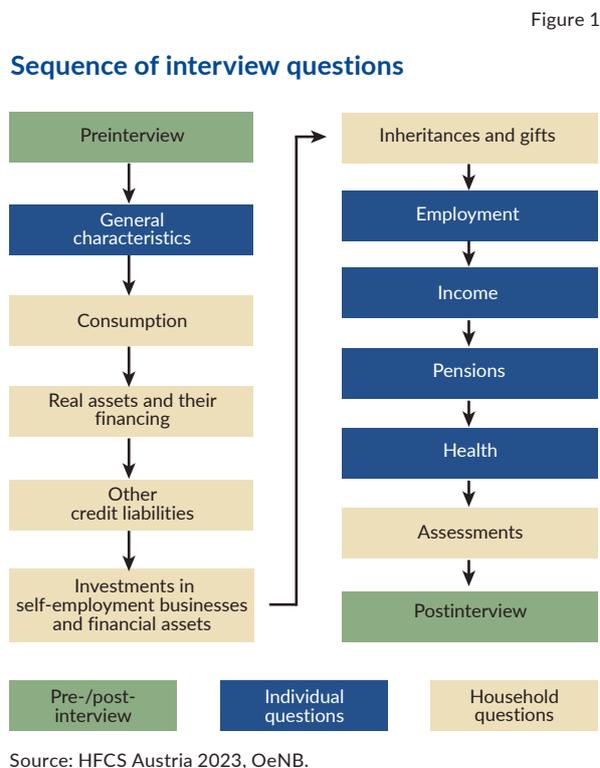## 2.4 Data collection period and reference period

In general, all the questions asked in the HFCS, especially those related to stock data, referred to the status quo at the time of the interview, which was conducted during the field phase from mid-March 2023 to early February 2024. In contrast, questions about income (aside from the question on the average monthly net household income)[4] referred to the calendar year 2022 (reference period). Of the interviews contained in the dataset, more than 85% were completed in 2021 and the remaining less than 15% at the end of the field phase in 2024.

Also for households interviewed in 2024, the reference period for yearly income was kept as 2022.

2.5 Interview structure and content

2.5.1 Questionnaire structure

The questionnaire consists of three parts: preinterview, main interview (divided into household and individual questions) and postinterview. This structure was chosen to make the survey as user-friendly as possible and to keep the duration of the interviews short. Figure 1 shows the sequence of interview questions in the HFCS.

Figure 1

Sequence of interview questions

Source: HFCS Austria 2023, OeNB.

---

[4] This is a noncore variable specific to Austria that is not included in the international HFCS dataset.

Preinterview

Before the actual interview, households were informed about the content and structure of the survey.[5] If a household was willing to participate, the interviewer first recorded the household matrix and identified the FKP. Recording the household matrix data involved determining the size of the household as defined for HFCS purposes, listing the individual household members and identifying the main respondent, i.e. the FKP or, if the FKP was not a household member, a reference person. It was also at this stage that key identifying data were collected for all household members, namely gender, age and relation to the main respondent. Finally, all these basic household data were listed in a table to facilitate verification and, if necessary, revision. If the FKP was present, the interview proceeded immediately – or, if the household so wished, at a later date – with the questions about general household characteristics.

General characteristics

In this section of the questionnaire, more detailed sociodemographic characteristics were collected for all household members: country of birth as well as length of stay in Austria for persons not born in Austria. For household members aged 16 and over, information about the level of education (including that of their parents) and marital status was recorded as well. The following sections (on consumption, real assets and their financing, other liabilities and credit constraints, investments in self-employment businesses and financial assets, inheritances and gifts) were used to collect information about the household as a whole.

Consumption

In this section, respondents were asked about the household's consumption and saving behavior and some components of household income. The questions on consumption aim at collecting information about households' typical average monthly expenditure on food, utility costs, other consumption products and services, total consumption spending and transfers to people outside the household. The monthly net household income was asked as a means to verify whether it was high enough to finance the household's expenditure; if not, further information was collected on how a household financed any expenses in excess of this income.

Expenses for holiday and travel were asked for the last calendar year. The questions about income and consumption were followed by topics about saving behavior and attitude.

Additionally, this chapter contained a set of questions concerning price expectations of households as well as information allowing the estimation of individual discount factors and some information on saving purpose.

Real assets and their financing

This section contained questions about the household's housing situation and most other real assets (excluding investments in self-employment businesses, which are addressed later) as well as their financing. The first set of questions established the location and size of a household's main residence and then focused on the tenure status of the household main residence (variable (A)HB0300), grouping households into (partial) owners, tenants or free users of their homes.

Households owning their main residence were asked to indicate when and how they had acquired it as well as the value of the property at the time of the interview and at the time they first acquired ownership.

---

[5] See chapter 3 for a detailed description of the contact strategy.

Further, homeowners were asked whether the property was being used as collateral for a loan. In the case of a collateralized loan, the following data were collected separately for a maximum of three loans: purpose of the loan, information on refinancing, initial amount, the year the loan was taken out, duration of the loan, outstanding principal, interest rate and type of interest, repayment rates and other characteristics. If a household had taken out more than three mortgages, the FKP was asked to provide summary information on the outstanding amounts and repayment rates of these additional (more than three) mortgages.

Households living in a rented home were asked about their rental cost (including and excluding running costs). Tenants of housing association apartments were asked to indicate the deposit made to the housing association. Any debt tenant households had incurred to finance such deposits was recorded under unsecured loans (see "Other liabilities and credit constraints").

Households that enjoyed free use of their main residence did not have to answer any further questions in the first half of this section.

The next set of questions for all three groups addressed additional real estate holdings, including holdings abroad. Specifically, loop questions, i.e. sets of questions with multiple iterations,[6] were used to identify basic information (i.e. type, size, value at the time of ownership transfer and at the time of the interview, and use) for up to three other real estate properties, as well as a summary question serving to establish the total value for any further properties. As was done with mortgages on the main residence, mortgages taken out on these properties were queried using loop questions for a maximum of three loans, and summary information was collected for further mortgages taken out on additional properties. In contrast to previous developments with regard to the ECB blueprint questionnaire, the implementation of this part on other real estate remained unchanged in Austria from previous waves to wave five.

Finally, households were asked to indicate the value of any cars or other vehicles as well as of any valuables (e.g. works of art, antiques) they owned. To conclude this section, respondents were asked whether they had bought a car or any other vehicles in the previous 12 months and if so, at what price.

Other liabilities and credit constraints

This category covered all other credit liabilities: leasing contracts, outstanding balances on current accounts and credit cards, private loans and noncollateralized loans.

For the first three types of liabilities, interviewers asked whether a household held such liabilities and how much was outstanding (for leasing contracts: the amount of lease payments per month). For outstanding balances on credit cards, the question put to households was: after paying the most recent monthly bill, was there any balance outstanding on your credit cards?

Respondents were asked separately whether they had private and noncollateralized loans. Private loans refer to loans from friends or family and noncollateralized loans to any further loans such as for example consumer loans or employer loans. When held by a household, the same information as for collateralized loans was gathered in the form of loop questions for the first three such loans of each of the two types. If a household had more than three of one of these two types of outstanding loans, the outstanding principal and the total monthly repayment rates for these were asked for in a summary question at the end.

This section of the questionnaire also addressed any loan requests that a household had made more recently and potential repayment difficulties of loans and other bills. Finally, respondents were asked about their attitudes to loans, their planning preferences and their level of risk aversion.

---

[6] See section 2.6.2 on the structure and navigation of loops.

Private businesses and financial assets

In the sections outlined in the following, the HFCS survey documented all wealth components of the household balance sheet beyond the real assets already covered. Interviewers started out by asking whether a household partly or entirely owned any nontraded self-employment businesses in which at least one household member played an active role, and, if so, recorded separate information for up to three such businesses (industry, legal form, number of employees and current value). If a household was invested and involved in any additional private businesses (other than the three business operations already recorded), the value of those holdings was asked at the end in a summary question. Another two questions related to ownership of any nontraded shares of businesses in which the household played no active role and to the value of these holdings.

The next questions focused on assets held in sight accounts, savings accounts, savings plans with building and loan associations, life insurance policies, mutual funds, bonds, listed stock, as well as assets in private foundations and managed accounts. For each of these items, FKPs were asked to indicate whether their household held such assets (yes/no question) and, if the answer was yes, what the total value of these assets was. For life insurance contracts, more detailed questions were asked. These included the date of conclusion, the type and duration of contract (benefits to be provided at the death of the policy holder or at a given date, or a combination thereof) and the frequency and amount of payments into the life insurance plan. This information allows for projections of the amounts held in such funds. Also, the fifth wave of the HFCS in Austria included several questions on the ownership, type and value of crypto assets. In addition, respondents were asked about money owed to the household as well as other financial assets. Interviewers also asked about a household's estimated total net wealth as well as for the distribution of this wealth across household members. This estimate was used to assess the plausibility[7] of the information provided, i.e. the sum of itemized figures was cross-checked with the total.

Inheritances and gifts

The next section of the questionnaire focused on the transfer of asset ownership in the form of inheritances and gifts. In a loop, interviewers asked about up to five inheritances or gifts,[8] recording information about the value at the time of transfer, the type and source of inheritance or gift, and when the transfer was made. The inheritances or gifts were listed in descending order, starting with the most important one for a household's current wealth. If there were more than five inheritances and gifts, their total value was queried in the questionnaire. However, in this survey for Austria (fifth wave) only three households indicated having received more than five.

Up to this point, almost all questions (of course, excluding those about the sociodemographic makeup of the household) focused on the household as a whole. In contrast, the questions in the three subsequent sections of the HFCS questionnaire, with a few exceptions, related to individual household members aged 16 years and above rather than to the household as a whole.

---

[7] For further details on consistency checks, see chapter 4.

[8] The international core dataset only contains up to three inheritances/gifts. In other words, it does not contain any additional inheritances/gifts. Over one-quarter of households reported intergenerational transfers and gifts. Of these below than 1% reported more than three.

Employment

Household members aged 16 and above were asked to provide information on their employment status. The first set of questions focused on people active in the labor market. Pensioners, homemakers, school-children, students and unemployed people were only asked about their expected retirement age and number of years spent in employment so far, before moving on directly to the second set of questions. Employed individuals answered questions about their occupation (ISCO code[9]), the number of working hours (with seasonal fluctuations), the employer's main economic activity (ÖNACE code[10]), the amount of time they had worked both for their current employer and overall in their working life, and about their expected retirement age. In the second part, all individuals aged 16 and above were asked to answer questions about their job history and their personal background.[11] The information provided in this section of the questionnaire is especially relevant in combination with that of the next section, which deals with income.

Income

Information on income was recorded by types of income. Respondents first indicated whether they received a certain type of income, and if so, what the annual amount was (information on annual income being usually readily available, e.g. on the income tax declaration). All types of income were recorded for the calendar year 2022 regardless of whether the interviews were completed in 2023 or 2024. Theoretically, for interviews that were conducted during 2023 it could be possible for a respondent to not yet know what their income in 2022 would have been, but for the interviews contained in this dataset this did not present a real problem. The option to add verbatim comments to responses made it possible to ensure during the editing phase that the information provided was indeed the information that had been queried (see chapter 4).

The types of income covered were employee and self-employment income, income from the state pension system and from private and company pension plans as well as income from unemployment benefits. For the first four types of income, respondents could provide either gross or net figures. (In the editing stage (see chapter 4) all net figures were converted to gross income figures using the Austrian Finance Ministry's gross-to-net calculator.)

Pensioners who indicated not to have received a pension in 2022 were queried about potential income in 2023. This information could be used to estimate 2022 income.

In addition to these individual questions, the household-level part of the interview with FKPs also established whether a household received income from regular (public or private) social transfers, from renting or leasing out real estate or from financial assets or private businesses. If respondents were unable to provide the gross income from financial investments, the relevant net income was also accepted. Finally, interviewers asked about other sources of income and about expected income growth.

---

[9] ISCO: International Standard Classification of Occupations, see https://www.ilo.org/publications/international-standard-classification-occupations-2008-isco-08-structure (accessed on July 17, 2025).

[10] ÖNACE: ÖNACE is the national version of the Statistical Classification of Economic Activities in the European Community (NACE – Nomenclature statistique des activités économiques dans la Communauté européenne); see https://www.statistik.at/datenbanken/klassifikationsdatenbank (accessed on July 17, 2025).

[11] The questions from the second part of this chapter in the questionnaire are not in their entirety part of the international core dataset.

Pensions

The main pension variables collected in the HFCS questionnaire included eligibility for future income from the state pension system and the number of state pensions to which respondents are entitled. Likewise, they were asked to indicate the number of contribution years and the account balances of, or their total contributions to, company or private pension plans.

Health

This section covers information about the household's FKP at the personal level. The focus is on this household member's health situation. The section includes questions about, e.g., the number of visits to a doctor, both financial and work-related information on health issues as well as questions related to general physical and mental health. Questions on long-term care insurance are also part of this section. It only contains country-specific information that is not comparable to information from other HFCS countries.

Assessments[12]

In this short section, FKPs were invited to provide their assessment of their households' position in the national wealth distribution as well as their view on the functions of wealth.

Postinterview

After the interview, respondents were encouraged to comment on questions they had found particularly hard to answer, items not covered by the questionnaire they would have deemed relevant, etc. Any comments were recorded as verbatim text. In addition, a so-called paradata section collected background information from interviewers about the interviews (see section 2.6.4.2).

2.5.2 Field phase

2.5.2.1 CAPI implementation (questionnaire programming)

The questionnaire was programmed using Warp-It developed by Solve-x software on the basis of a Word template, the German original of which is available in the online appendix. In addition to filtering, the questionnaire also prompts internal consistency checks (see chapter 4) that will flag possible errors in data entry during the interview. The use of the computer-assisted personal interviewing (CAPI) technique combines the advantages of a personal interview with those of real-time digital recording and data crosschecking. In addition, it allows for the implementation of complex filtering techniques, producing tailormade questionnaires for each household.

2.5.2.2 CAPI test

After the first programming phase, the questionnaire was tested by members of the OeNB's HFCS team and in a pilot survey of 50 households.

---

[12] This is a noncore variable specific to Austria that is not included in the international HFCS dataset.

2.5.2.3 CAPI problems

This wave of the HFCS benefited from the experiences of the four prior waves. As a result, errors could be largely eliminated. Below is a list of programming issues that led to reprogramming during the field phase of the fifth wave or editing thereafter:

- Acquisition crypto assets (AHD1770): At the start of the field period, the other category verbatim recording was shown twice. Due to a quick fix, such cases had no negative implications for the interview.
- Check the year of taking out unsecured loans ((A)HC059$): For the first interviews, the consistency check of this variable excluded 2023 as a potential answer. No information was lost and the consistency check was updated early on.
- Net income from financial assets (AHG0420): In the first weeks of the field period, the reference to the interval list in case of net income with fixed interval answer was wrong and had to be corrected.
- Maximum number of digits (few variables: AHC100$, APF0930, APF0940): At the start of the field period, only 6-digit recording was possible while this was extended to 9-digit numbers. No affected household could be detected.
- Likelihood to lose current job (pez010): Filtering with regard to occupational status ((A) PE00200x) wrongly regarded only main status instead of main and secondary status. No appropriate information of household members was lost before correcting this flaw.

2.6 Special features

2.6.1 Loops

Various aspects of a household that are especially relevant for the HFCS were asked using loops, i.e. sets of identical questions to collect information on individual items applicable to a household (e.g. loans). Figure 2 shows the schematic cycle of these loops.

Loops were used for the following items:
- mortgage loans using the household's main residence as collateral (max. number of iterations: 3),
- real estate holdings other than the household's main residence (max. number of iterations: 3),

Figure 2

Structure of a loop

Source: HFCS Austria 2023, OeNB.

- mortgage loans that used these other properties as collateral (max. number of iterations: 5),
- unsecured loans from family and friends (max. number of iterations: 3),
- other unsecured loans (max. number of iterations: 3),
- investments in self-employment businesses (other than holdings of listed shares) (max. number of iterations: 3),
- life insurance contracts (no max. number of iterations),
- ownership transfers through inheritances or gifts (max. number of iterations: 5).

First, respondents indicated whether a specific item applied to their household, and if so, the number of such items held by the household. The details were then recorded in a loop for each of these items to the maximum number of iterations. For instance, if a household held two other unsecured loans, the respondent was asked to first provide details of the loan with the higher outstanding balance and then details of the second loan. If a household had received multiple inheritances and gifts, the procedure was repeated up to five times. When there were more items than the maximum number of iterations permitted in the loop, interviewers recorded a summary of the information.[13] For life insurance contracts, the questionnaire did not set a maximum number of iterations in order to be able to record all contracts individually.

To make the survey as user-friendly as possible, it was made possible for interviewers to exit a loop at any time. In this case, information was collected in summary questions. Interviewers were instructed not to overuse the option of exiting loops; this feature was only meant to prevent respondents from breaking off the interview early if they were unwilling to answer a given question for more than one item.

The loops for real estate holdings other than the household's main residence and mortgage loans collateralized with these other properties are structured differently in the ECB template questionnaire. Whereas the Austrian questionnaire used two separate loops to establish this information (as in wave one and wave two), the questions on loans collateralized with further property were nested in the loop about further property in the ECB blueprint questionnaire. Additionally, the HFCS in Austria was designed to record in this section which property was being used as collateral for a particular loan. Thus, all the information required by the ECB questionnaire was contained in the Austrian questionnaire and the variables were adapted and renamed as required for the ECB version after the data imputation phase.

2.6.2 Euro loops

All questions involving monetary amounts were asked in a loop to avoid data entry errors and to obtain a range containing the approximate amount if respondents were unable or unwilling to state specific amounts. This section describes the structure of euro loops (see the online appendix, "Euro loops," for a schematic overview).

In the first step, respondents were asked to provide specific amounts ("How much…?") in any given currency. After the information had been recorded, respondents were asked to confirm the amount and the currency. ("You said the amount is … [currency]. Is that correct?")

If no specific amount was provided, respondents were asked to indicate a range ("Could you provide a range for the amount?"). Rather than having an upper and a lower bound, the range could also be limited at one end and open at the other end (e.g. "no more than EUR …" or "at least ATS …"). If respondents stated a range, the interviewer continued again by asking respondents to first indicate the currency and then confirm the range and the currency.

---

[13] There were very few cases where respondents reported a higher number of items than the maximum number of iterations in the loop.

Table 1

List of predefined ranges for euro questions

	List A		List B		List C
EUR					
A	1 – below 101	A	1 – below 10,001	A	1 – below 1,001
B	101 – below 201	B	10,001 – below 50,001	B	1,001 – below 2,501
C	201 – below 301	C	50,001 – below 75,001	C	2,501 – below 5,001
D	301 – below 401	D	75,001 – below 100,001	D	5,001 – below 7,501
E	401 – below 501	E	100,001 – below 150,001	E	7,501 – below 10,001
F	501 – below 751	F	150,001 – below 200,001	F	10,001 – below 15,001
G	751 – below 1.001	G	200,001 – below 300,001	G	15,001 – below 20,001
H	1,001 – below 1,501	H	300,001 – below 400,001	H	20,001 – below 25,001
I	1,501 – below 2,001	I	400,001 – below 500,001	I	25,001 – below 30,001
J	2,001 – below 3,001	J	500,001 – below 750,001	J	30,001 – below 35,001
K	3,001 – below 5,001	K	750,001 – below 1 mio	K	35,001 – below 40,001
L	5,001 – below 7,501	L	more than 1 million – 3 million	L	40,001 – below 50,001
M	7,501 – below 10,001	M	more than 3 million – 5 million	M	50,001 – below 75,001
N	10,001 – below 25,001	N	more than 5 million – 10 million	N	75,001 – below 100,001
O	25,001 – below 50,001	O	more than 10 million	O	100,001 – below 200,001
P	more than 50,001			P	200,001 – below 300,001
				Q	300,001 – below 500,001
				R	500,001 – 1 million
				S	more than 1 million

Source: HFCS Austria 2023, OeNB.

If respondents were unable ("Don't know") or unwilling ("No answer") to indicate an individual range themselves, they were asked to choose a range from a list. Depending on the question at hand, interviewers used one of three lists of ranges (see table 1).

The three lists of predefined ranges (A to C) are based on the (unweighted) empirical distribution of the answers to numerous questions in the previous waves of the HFCS in Austria. This evidence showed that, for specific questions, the main part of the distribution called for smaller and hence more specific ranges than the remaining parts of the distribution. List A was used for questions about consumption expenditure and loan repayments. List B was used for questions related to properties and investment in self-employment businesses, and list C was typically used for outstanding loans and incomes. Questions about financial assets were aligned either with list A or list C, depending on the distribution of assets as observed in the first wave of the survey.[14] The predefined ranges referred to amounts in euro only. Interviewers moved on to the confirmation question as soon as the respondent had chosen a range. The lists of predefined ranges were presented in the form of showcards for all questions involving monetary amounts (see also section 3.5.4).

Only if a respondent also refused to choose from the list of predefined ranges was the status of the question recorded as not answered ("Don't know" or "No answer"). The information recorded with the ranges was especially important for multiple imputation (see chapter 5).

---

[14] See the questionnaire in the online appendix at https://hfcs.at/en/results-tables.html for a detailed overview of which ranges (list A, B or C) were used for which questions.

2.6.3 Recording farming households

Recording the (real) assets of farming households was found to be a particular challenge for respondents in the first waves, especially when it came to breaking down the assets into the household main residence and business assets. The business assets of farmers have been recorded in the loop on investment in self-employment businesses since the first wave. In order to elicit more precise answers, the second-wave questionnaire introduced a number of additional questions for farming households as well as some additional guidance. These additions and enhancements were retained in the questionnaire for the subsequent waves. The procedure can be summarized as follows:

- Before the actual interview started, respondents were classified by interviewers as running an "agricultural business" or as running "no agricultural business." The classification was straightforward in all but a few cases. But even in the few cases where a respondent was incorrectly classified, the structure of the questionnaire ensured that all essential information was still obtained.
- Specifically, the following extra information was recorded for households classified as farmers:
  - Was it possible to separate housing assets (i.e. the household main residence) from business assets? (in the "main residence" chapter of the questionnaire)
  - If not, what percentage of the recorded value did respondents allocate to their main residence? (in the "main residence" chapter of the questionnaire)
  - Does the value recorded for investment in a self-employment business include the main residence recorded? (in the "investments into self-employment businesses" chapter of the questionnaire)
- Moreover, for questions relating to the value of their main residence, the yes/no question on properties other than the main residence, as well as for the question about investment in a self-employment business and its value, farming households received detailed guidance as to which components of their household balance sheet were to be recorded under which position.

In addition, all interviewers were specifically trained to handle such cases (see also section 3.3). The additional information thus collected proved to be particularly relevant for multiple imputation (see also section 5.4).

2.6.4 Other information recorded

2.6.4.1 Contact attempts

Only after unsuccessful attempts at contacting a household on at least five separate occasions could a household be classified as a unit nonresponse (see also the contact rules in section 3.4).[15]

The contact attempts were recorded in the dataset,[16] thus providing additional information. The exact time (year, month, day, hour and minute), mode and outcome of every single contact attempt were documented, as was the total number of contact attempts. Interviewers were instructed to write this information down on paper first and record it in the electronic questionnaire only following the completion of the work (interview, refusal, etc.) on a particular household.

---

[15] Interviewers had to make at least two contact attempts in person over a period of at least three weeks.

[16] These variables were not included in the user database due to anonymization requirements.

2.6.4.2 Paradata

Two kinds of so-called paradata were collected: While the first type of paradata was collected for all households – including those that ultimately did not participate in the survey – the second type covered additional information on the households that were interviewed.

The first section covered all information that could be obtained without actually entering a household's residence or completing an interview: the interviewer's assessment of the building and construction type, the geographical location (urban or rural area), the condition of the building, the residential area and special security measures.

If an interview took place, interviewers also collected the following additional paradata: the condition of the dwelling's interior, the interview language (in Austria all interviews were conducted in German), the interviewer's assessment of the accuracy of the information provided, the village or town in which the interview was conducted, the number of people present during the interview and the interest they showed in the interview, the frequency with which respondents consulted documentation to answer questions, and the type of documentation used. In addition, interviewers had to submit written comments about the interview for every single household. These comments, to be made on five questions covering the interview as a whole, proved very helpful at different stages of the project.

The first section of paradata was recorded in the sample register file, which is not part of the user database due to anonymization requirements. It was used mainly to calculate nonresponse weights.[17] The second section (excluding interviewer comments) was recorded in variables HR0100 to HR1600 in the household data file, which is part of the HFCS dataset.

2.7 Interviewer documents

Among other things, interviewers had access to the following documentation to help them prepare for an interview and as a reference point during the interview (also available in the online appendix):[18]

- the showcards which were used during interviews to provide respondents with a list of response options for several questions in the questionnaire,
- a glossary which contained simple definitions of the terminology used in the questionnaire, and
- a copy of the study entitled "Vermögen der privaten Haushalte in Österreich – Gemeinsamkeiten und Unterschiede" (Fessler and Schürz, 2019) to illustrate how the data obtained in the third wave of the HFCS in Austria were used for analysis.

---

[17] These weights are used for the correction of the nonrandom participation of households in a survey and are needed to construct the final household weights (see chapter 7).

[18] For a detailed description of the documents, see chapter 3.

2.8 Participating countries

The fifth wave of the Eurosystem HFCS was conducted in the following euro area countries: Belgium,[19] Croatia, Germany,[20] Estonia,[21] Ireland,[22] Greece, Spain,[23] France,[24] Italy,[25] Cyprus, Luxembourg,[26] Latvia, Lithuania, Malta,[27] Austria, the Netherlands, Portugal, Slovenia,[28] Slovakia,[29] and Finland.[30] Additionally, the fifth wave of the HFCS is expected to include Czechia and Hungary.

The survey was prepared by the Household Finance and Consumption Network (HFCN) launched by the ECB. The aim was to achieve ex ante harmonization at as many levels of the survey as possible. In doing so, it was necessary to take features specific to individual countries into account, leading to discrepancies and to additional questions in some cases. In addition to the core output variables, the Austrian survey also collected data that are specific to Austrian households (e.g. information on foreign currency loans). Moreover, the answer options for some questions were categorized in greater detail in the national datasets. A case in point is the question about respondents' marital status, which came with six answer options in the national dataset but only five answer options in the international dataset. The OeNB is planning to provide the country-specific details as additional information to the datasets that the ECB is expected to publish in spring 2026.

2.9 Online appendix

The following PDF documents are available in German for download from the Austrian HFCS website (https://www.hfcs.at/en/results-tables/) as an appendix to this chapter:

- the questionnaire,
- the euro loops,
- the paradata questions,
- the variable lists,
- the showcards, and
- the glossary.

[19] Information on the survey in Belgium is available at https://www.nbb.be/en/publications-and-research/publicatip ons/other-publications/household-finance-and-consumption-survey-0 (accessed on August 11, 2025).

[20] Information on the survey in Germany is available at https://www.bundesbank.de/en/bundesbank/research/paa nel-on-household-finances-phf (accessed on August 11, 2025).

[21] Information on the survey in Estonia can be found at https://www.eestipank.ee/en/statistics/research-financial-behaviour-and-consumption-habits-estonian-households (accessed on August 11, 2025).

[22] Information on the survey in Ireland is available at https://www.cso.ie/en/methods/socialconditions/hfcsurvey/ (accessed on August 11, 2025).

[23] Information on the survey in Spain is available at https://app.bde.es/efs_www/home?lang=ES (accessed on August 11, 2025).

[24] Information on the survey for France, where it is run under the Enquête Patrimoine, can be found at https://www.insee.fr/fr/metadonnees/source/serie/s1005 (accessed on August 11, 2025).

[25] Information on the survey in Italy is available at https://www.bancaditalia.it/statistiche/tematiche/indagini-famigb lie-imprese/bilanci-famiglie/index.html?com.dotmarketing.htmlpage.language=1 (accessed on August 11, 2025).

[26] Information on the survey in Luxembourg is available at https://www.bcl.lu/en/Research/enquetes/hfcs/about/ index.html (accessed on August 11, 2025).

[27] Information on the survey in Malta is available at https://www.centralbankmalta.org/household-finance-and-con/ sumption-survey (accessed on August 11, 2025)

[28] Information on the survey in Slovenia is available at https://www.bsi.si/en/household-finance-and-consumption-survey (accessed on August 11, 2025).

[29] Information on the survey in Slovakia is available at https://nbs.sk/en/research-at-nbs/household-finance-and-consumption-survey-hfcs/ (accessed on August 11, 2025).

[30] Information on the survey in Finland can be found at https://stat.fi/en/statistics/vtutk (accessed on August 11, 2025).

3 Interviewers

This chapter provides an overview of the HFCS interviewers' role and tasks. It describes how interviewers were supported and monitored in their work and how the data they collected were examined.

3.1 The interviewers' role in the survey process

The information on households collected in the HFCS in Austria is generally considered to be sensitive. Therefore, the personal interviews conducted by trained interviewers played a major role in the survey process. Interviewers' professionalism, profound knowledge of the survey's subject matter, excellent interviewing skills and appropriate behavior are a precondition for surveys to be successful and therefore contribute in particular to the quality of the resulting data. To prepare for the HFCS, interviewers completed comprehensive training on the content and structure of the HFCS.

In the field phase and during the personal interviews, it was possible for interviewers to consult written reference material and, if necessary, receive support from the OeNB.

3.2 General information

The number of interviewers involved in the fifth wave of the HFCS was 48. While the survey company decided which interviewers to involve in this complex and sensitive survey, the OeNB reserved the right to withdraw individual interviewers if they did not meet the quality criteria.

In general, the interviewers had specific experience conducting household surveys, having been involved either in past waves of the HFCS in Austria or in surveys of a similar magnitude (e.g. the OeNB Household Survey on Housing Wealth 2008, EU-SILC or SHARE). In fact, almost 90% of the interviewers in the fifth wave of the HFCS in Austria had also conducted interviews at least in one of the previous waves. Payment for successfully completed interviews was calculated on the basis of the surveyed household size; a considerably lower remuneration was paid for the collection of paradata when interviews were not completed successfully. Travel expenses were also refunded. To be entitled to a refund of travel expenses for uncompleted interviews, the interviewers were required to have made at least two personal contact attempts and five contact attempts altogether. Particularly efficient interviewers received bonus payments.

3.3 Interviewer training

All interviewers conducting interviews in the HFCS were specially trained. The training content was developed by the OeNB in cooperation with the survey company. The development and intention of the interviewer training was based on the experience of previous waves. All training sessions were conceived as full one-day trainings. Due to positive experiences with online trainings during the COVID-19 period, we also conducted the sessions as digital meetings. In order to keep participants' concentration high and the digital workshops interactive, trainings were split into two half-day gatherings. The contents remained unchanged and are explained below. Additionally, interviewers had to conduct a take-home interview by themselves. Questions arising in this take-home exercise could be answered in the second half of the digital training sessions.

There were six distinct training sessions in total. The first one was a training slot for particularly experienced interviewers that also conducted the pre-test interviews.

3.3.1 Training unit 1

Introduction

First, a member of the OeNB HFCS team introduced the interviewers to the topic and the aims of the HFCS in Austria. This introduction also covered information about the use of data, including explanations why a central bank requires the data surveyed and how researchers use the data and communicate results to the media. Knowledge of these issues is considered to help interviewers' motivation. The HFCS team representative also described the use of data and analytical approaches on the basis of examples and emphasized the importance of conducting interviews conscientiously and of all households in the sample taking part in the survey. Finally, the central role of interviewers in the HFCS data collection process was highlighted.

Overview of the questionnaire

Following the introduction, the participating interviewers were made familiar with the questionnaire: Its chapter structure, the definition of "household" within the meaning of the HFCS, the identification of financially knowledgeable persons (FKP), how to distinguish an FKP from a reference person, loops, the method used for recording amounts in euro (including the structure of a (euro) loop, see section 2.6.2).

3.3.2 Training unit 2

The briefing on the questionnaire started with a theoretical introduction, supported by additional information and documentation, where required. After that, the lecturer walked the workshop participants through the CAPI questionnaire using an unrealistically complex household as an example. This approach made it possible for participants to acquaint themselves with the essential elements of the questionnaire both in theory and practice. This training unit was split into the two blocks described below.

Questionnaire – theory

The first part of this training unit covered the preinterview questionnaire including the creation of a household matrix and the selection of the household's financially knowledgeable person. In addition, the general characteristics of the household members, the questionnaire section on the household's consumption behavior and the household's real estate wealth and its financing were discussed. Explanations of how to treat farming households were also given ample time.

The treatment of other liabilities, private businesses and financial assets, as well as the section on inheritances and gifts were introduced. In particular, participants were walked through the range of financial assets to address possible misunderstandings, and they learned about the fundamentals of the stock and flow data in households' balance sheets and how to record additional comments.

Lastly, the training provided information on household members' employment status, income and retirement provisions. The training moved toward incomes at the household level and assessment questions. In particular, participants were acquainted with the reference period for income as well as the options for recording income (gross or – if the gross amount was not known – net of tax and social security contributions). The lecturers highlighted the importance of recording comments provided by the respondents.

Questionnaire – practice

Part two of this training unit covered the simulated interview. As in the previous waves, interviewers asked the questions, and the workshop leader gave the answers. All types of questions were covered in this

mock interview. This means that not all the questions could be reached but in order to understand and discuss a question not all the iterations have to be gone through. So the interview covered more than one item of one loop and one mortgage, for example, to explain all the details of a loop and of loans, but it was not necessary to cover three iterations for every possible type of loan. This part of the workshop was particularly interactive. Interviewers asked ample questions that could be addressed and resolved.

3.3.3 Training unit 3

Interviewers' tasks, contact specifications and paradata

The key initial task for interviewers was to convince the selected households to take part in the HFCS. In this respect, the interviewers were provided with a comprehensive list of reasons in favor of participating, as well as information on data security and the contact details for people at the survey company and the OeNB who they could turn to in case of problems. In training unit 3, the interviewers were given exact, detailed specifications on how to proceed when contacting households (see section 3.4). Among other things, interviewers were required to document their attempts to contact the selected households and compile all paradata (see section 2.6.4). The lecturers highlighted in particular that accuracy in compiling information was of utmost importance and that interviewers thus contributed substantially to data quality.

Guidance on interviewer communication

In the second part of this training unit, interviewers received guidance on how to communicate during interviews, for instance with regard to providing explanations or querying answers. In addition, they were trained not to express their personal opinions if respondents asked them questions. Likewise, interviewers learned to repeat and explain questions in the most neutral way possible (using the glossary, if necessary). Comments received from the previous waves of the HFCS in Austria helped to highlight typical interview situations.

3.3.4 Training unit 4

Documents and other material

In training unit 4, lecturers and interviewers once again went through all the documentation and material made available to the interviewers, which had been used in training units 1 to 3 (see section 3.5). This provided the participating interviewers with another opportunity to ask questions on all aspects of the HFCS.

Organizational information

Finally, interviewers were provided with organizational information, such as the addresses of households that they had to contact. Also, they received information about the incentives for households that completed an interview and interviewers' remuneration.

3.4 Contact strategies and specifications

The process of establishing contact with the households in the HFCS sample took place according to detailed specifications provided by the OeNB. One or two weeks prior to the first contact attempt by the interviewer, the survey company sent the households selected in the sample an individualized advance letter signed by the OeNB Governor as well as an information leaflet. This prior notification enabled

respondents to prepare in advance for interviewer visits. By consulting the information material provided, as well as the HFCS website (www.hfcs.at/en), households were able to familiarize themselves with the survey topic, consider whether they wanted to take part and, if so, prepare useful documents (such as bank account statements, etc.).

With the advance letters having been sent, interviewers had to make up to five contact attempts with each household. At least two of these contact attempts were to be made personally (by visiting the household's address in person and trying to establish contact); at least one attempt was to be made at the weekend and another outside normal working hours (9:00 a.m. to 5:00 p.m.). All contact attempts had to be spread out over a period no shorter than three weeks. This approach was necessary in order to rule out distortions as a result of selective participation (e.g. many single-person households cannot be reached during the day and can only be contacted in the evening or at the weekend).

The interviewers were required to document each contact attempt. During at least one of the personal contact attempts, information on the exterior and the location of the property (see also section 2.6.4.2 on paradata) was recorded, even if no successful interview took place with the household in question.

The interviewers were instructed to carry with them all the necessary material (notebook computer, information material, participation incentives, etc.) during each personal contact attempt. This allowed them to react appropriately to different situations, e.g. if a household wanted to participate in the survey immediately, if they requested time to consider or wanted to make an appointment, or if they declined to be interviewed. If requested, interviewers also had to offer interview appointments at the weekend or in the evening as well as the option of meeting respondents outside their main residence (e.g. at the respondent's office).

3.5 Documents and other supporting material

In addition to the specific training the interviewers received upfront, interviewers were provided with the following information and supporting material to be used during the interviews, where appropriate:

3.5.1 Letter by the OeNB Governor

Shortly before the first personal contact attempt, all households received an individualized letter and an information leaflet (see online appendix) explaining what the survey was about, what objective it served, who to contact in case of questions, how the collected data would be used and that all data would be treated confidentially. Interviewers took this letter, which was signed by the OeNB Governor, with them whenever they contacted households.

3.5.2 Incentives

As participation in the survey was voluntary, monetary incentives were used to increase households' willingness to take part in the HFCS. Each household that successfully completed an interview received a Sodexo voucher with a value of EUR 15. The interviewers handed over the voucher to the respondents directly upon completion of the questionnaire.

3.5.3 Scientific study

The interviewers were instructed to have with them a copy of the study "Vermögen der privaten Haushalte in Österreich - Gemeinsamkeiten und Unterschiede" by Fessler and Schürz (2019) (see online appendix) during each contact attempt. This study is based on data taken from the third HFCS wave in Austria and

serves as an example of how survey data are used in a statistical context. Respondents thus had the opportunity to inform themselves how the information they provided was going to be used, which helped to increase confidence in the survey. Interviewer feedback after the first and second waves showed that initially reluctant respondents were more likely to participate in the survey after having received this information.

3.5.4 Showcards

To answer certain questions of the survey, respondents had to choose from a list of answers presented by the interviewer on showcards (see online appendix), which covered the following topics:
- Euro amount ranges A
- Euro amount ranges B
- Euro amount ranges C
- Questions for capturing the demographics of household members
- Relation to the reference person
- Housing costs including service charges
- Expenditure for travel and holidays
- Expenditure for other consumer goods and services
- Types of income
- Unexpected windfall gain – lottery
- Five value changes
- Reasons for not renting out
- Economic sectors
- Types of life insurance contracts
- Types of mutual funds
- Investment behavior
- Type of inheritance/gift
- Employment status I and II
- Employment classifications
- List of possible probabilities
- Health
- Well-being
- Personality traits
- Functions

The questions that required interviewers to use a showcard were specifically marked in the questionnaire. The digital version of the questionnaire also contained references at places where the use of a showcard was required.

3.5.5 Contact form

Interviewers were required to document all information on contact attempts initially by hand on the contact form, which, upon conclusion of a household interview, was digitized with the same software that was used for the questionnaire.

Aside from the household's identification number, the documentation comprised the date, time, type (e.g. personal or by telephone) and outcome (e.g. complete interview or ineligible address) of a contact attempt. Personal identification information (such as name, address or telephone number) was not part of the data and was not forwarded to the OeNB.

3.5.6 Interviewer manual

The interviewer manual distributed to all interviewers included all necessary information on the HFCS (e.g. the definition of a household) and served in particular as a reference point for the interviewers. In addition to an introduction to the questionnaire, its special features (see chapter 2) and all related documents, the manual also outlined the tasks of the interviewer. Furthermore, it provided guidance on how to locate households and convince them to take part in the HFCS. It also described the requirements for interviewer behavior and their interaction with the people contacted. Other important features were detailed contact specifications and answers to questions frequently asked during the first contact attempt. The manual additionally comprised essential legal texts on data protection that the interviewers had to be familiar with. Furthermore, the manual listed the contact data of the survey company (including a hotline telephone number) and the telephone number of the OeNB hotline in case the interviewers had any questions. The interviewer manual provides an extensive overview of the preparations for the HFCS and can therefore be found in the online appendix.

3.5.7 Glossary

Working for the HFCS required a basic understanding of a broad range of different financial instruments, investment opportunities and types of income, as well as the acquisition of real assets. Interviewers had at their disposal an alphabetical glossary (see online appendix) that provided explanations of technical terms. The glossary consisted of some 20 pages of explanations for all terms of key importance to the HFCS, such as "mutual fund" or "household" (according to the HFCS definition).

Already at the training stage, the interviewers were instructed to use this glossary to acquire relevant knowledge which they would be able to fall back on during interviews. By virtue of its references to the variables recorded in the survey, the glossary is also of importance when analyzing the collected data, as it explains the technical terms contained in the questionnaire.

3.6 Monitoring

To uphold the high quality standards of the HFCS, both the survey company and the OeNB monitored interviewer performance. The interviewers' direct contact person and superior was a regional area manager who reported to field management at the central office in Vienna. The survey company monitored in particular the correct execution of the interviews by checking roughly one in every six interviews via telephone from Vienna. During these calls, the contacted respondents were asked to provide data on the composition of their household, the conduct and duration of the interview and the topics covered.

Furthermore, the data from completed household interviews were forwarded to the OeNB promptly, in 24 batches (including answers to queries) during the field phase, to enable OeNB staff experts to monitor interviewer performance in a timely manner (see section 4.4.1). In addition, the following interviewer performance indicators were examined: item nonresponse (both broken down by real assets and financial assets and in aggregate form for the entire interview), the relative duration[1] of an interview, the number of questions asked, the number of households surveyed successfully and unsuccessfully, and the resulting unit nonresponse, as well as the number and quality of interviewers' comments. The specific comments to be made by the interviewers upon completion of each household interview were also examined.

---

[1] During each interview, time logs were recorded at different points in the questionnaire.

The OeNB's goal in this phase was to quickly identify and resolve difficulties with prompt analysis. Monitoring interviewers gave the OeNB a chance to address individual interviewers' difficulties concerning certain topics or aspects by providing targeted guidance. The OeNB also had the possibility to withdraw interviewers with immediate effect from the survey if they did not meet the quality requirements.

Additionally, a high-level steering committee meeting with participants from the OeNB and the survey agency was held regularly (in general every 2-4 weeks) to address potential problems directly.

3.7 Problems relating to interviewers

Shortcomings identified during the monitoring process were pointed out to the interviewers. For instance, if interviewers had difficulties entering the correct number of zeros for (large) numbers – a problem that was relatively easy to identify with the help of the numerous plausibility checks – they were asked to pay particular attention in subsequent interviews. The next batch of data was then examined for the persistence of these problems. In the case of some of the interviewers, monitoring also helped reduce the item nonresponse rate of the households they interviewed.

No interviewer had to be withdrawn in the fifth HFCS wave in Austria.

3.8 Survey of interviewers

The HFCS in Austria also entailed the systematic collection of information on the interviewers involved. The information provided by the interviewers on a voluntary basis included socio-economic information (age, gender, education, region), employment status including work experience as an interviewer, personality-related indicators and the interviewers' financial situation. Interviewers also had the opportunity to document their experience working for the HFCS in Austria. This information is particularly relevant for the nonresponse adjustment of the complex survey weights (see chapter 7). The participation rate in the survey of interviewers was close to 100%.

3.9 Online appendix

The online appendix includes the letter by the OeNB Governor to the households, the information leaflet, the showcards, the interviewer manual, the alphabetical glossary as well as the exemplary study by Fessler and Schürz (2019).

4 Consistency checks and editing

4.1 Introduction

Here, data editing is understood as the amendment of electronically recorded observations collected through individual interviews, so as to correct any errors or logical inconsistencies that may have occurred during the survey, as well as the aggregation of information that was recorded via auxiliary variables, typically with a view to keeping the questionnaire as clear and user-friendly as possible. The editing process is thus essential for improving the quality and consistency of the datasets.[1]

The raw data collected in surveys do not always contain the information that the questions were intended to elicit. As respondents in the HFCS occasionally either experienced difficulties in understanding the questions asked or had insufficient knowledge on the substance of the survey, they may sometimes have provided inaccurate information. At the same time, data entry errors may have occurred (see also chapter 3), or data may have been processed inaccurately. In the HFCS, great importance was attached to minimizing such errors up front through the structure and wording of the questionnaire; through checks and, if needed, adaptations during the field phase; and ex post through data editing, as will be outlined in this chapter.

This chapter provides insights into the consistency analyses and edits performed for the fifth HFCS wave in Austria, starting with information on the number of edits performed (section 4.2) and followed by explanations on the consistency checks conducted during and after the interviews (sections 4.3 and 4.4). Furthermore, we outline the flags used to highlight ex post adjustments of the observations recorded (section 4.5), provide a detailed account of ex post editing (section 4.6) and describe formatting and editing after multiple imputations (section 4.7). The chapter ends with concluding remarks (section 4.8).

4.2 Number and type of edits

All in all, around 44,600 of the more than 1.4 million observations collected in the fifth HFCS wave were edited, i.e. about 3.1% of all data points were amended (see table 2). This figure is comparable to the last waves of the HFCS in Austria and hence expected.

The row "Total" indicates the full range of edits that were implemented. Edits that resulted in actual changes to the collected values, i.e. real changes, were limited to 6,423 observations (see row "Edits based on expert judgment and follow-up phone calls"), which corresponds to a change rate of 0.4% (slight reduction from previous waves). These changes involved primarily inconsistent values that were corrected based on subsequent queries and/or other information or were deleted and replaced through imputation. Just a little over 60% of all amendments (see row "Edits based on other survey information (e.g. verbatim records)", i.e. a little over 28,400 observations, could be derived from the verbatim records and the use of a flexible questionnaire design for certain questions (e.g. questions about life insurance policies or total annual net income). All in all, around 2% of all observations were amended through this type of editing, making them the most prominent form. This indicates how important it is to allow for verbatim records on a large scale. Questionnaires as detailed as the one used for the HFCS in Austria must be user-friendly to ensure the participation of respondents and high-quality standards. Various data – e.g. data on the occupation (ISCO code) of employed individuals – are only collected as verbatim responses to minimize the effort required from respondents. The flexibility for certain questions was achieved by giving respondents more options on how to respond to the question. For some income questions that asked for

---

[1] See e.g. Kennickell (2011) and Bledsoe and Fries (2002) for information on the editing measures used in the Federal Reserve's Survey of Consumer Finances.

Table 2

Number and type of edits

	Total observations[2]	Number of edits	Share of edited observations in total observations
			%
Total[1]	1,436,775	44,609	3.1
Edits based on expert judgment and follow-up phone calls		6,423	0.4
Edits based on other survey information (e.g. verbatim records)		28,441	2.0
Deleted observations		9,745	0.7

Source: HFCS Austria 2023, OeNB.

[1] The line "Total" contains all edits.
[2] Includes only observable information. Filter missings are excluded.

the gross value, net values could be provided instead if respondents did not know the gross value. The net value was then converted to its gross value during editing.

In around 9,700 cases (i.e. around 0.7% of observations), observations were set to missing (".").[2] Edits of this type were made for different reasons, but mostly in the process of data cleanup (see section 4.6.2.1). Further, some items had been entered in a wrong position in the questionnaire. When transferring such information to the right position, the original entry had to be deleted, i.e. set to missing. In some cases, complete sets of entries were set to missing ("."), among other things, because the corresponding head variable had been edited. Below is an example[3] that demonstrates the editing process resulting from some of the reasons just outlined.

A case in point would be the duplicate recording of income from pensions, first under "Received employee income" and then under "Received income from public pensions." Here, the head variable "Received employee income" (PG0100) was changed to "No" and the value recorded for this variable was deleted because the respective income figure had been adequately recorded under the pension income variable (PG0300 and PG0310).

4.3 Consistency checks during interviews

The HFCS is based on computer-assisted personal interviewing (CAPI). CAPI has a number of advantages over the use of paper-based questionnaires or phone-based interviews. Interviewers walk the respondents through the questionnaire on screen, using a laptop on which the survey software is installed. The information collected is checked for integrity and consistency as it is being entered. Any questions of clarification that the respondents may have raised can be resolved immediately either by the interviewer or with the aid of the supporting documentation, and thus errors can be prevented during data entry.

However, consistency checks during an interview are subject to limitations in terms of scope. An excessive number of consistency checks during an interview would make it exceedingly long and thus wear out the respondents and in turn decrease the standard of the data collected and/or might even cause respondents to break off an interview.

[2] The cleanup statistics do not reflect irrelevant variables cleaned up following the skipping of certain questions in a loop (see also sections 2.6.1 and 4.6.2.4).

[3] Examples given in this chapter are in italics.

Moreover, restrictions arise from the fact that all information which should be used for the consistency checks must already be available. These limitations do not apply to simple consistency checks linked to specific predefined benchmarks. Whenever certain limits are exceeded or undercut, pop-up warnings appear that allow the entry to be checked immediately. However, the information necessary for more complex consistency checks often does not become available until answers are received in later stages of the interview.

The digital version of the questionnaire used for the HFCS provided for close to 250 consistency checks,[4] typically in the form of "soft" checks that highlight potential issues but do not prevent the interview from proceeding if the interviewer decides to continue. Whenever a test criterion was violated, a warning message popped up.

For example, if a household with a disposable net monthly income of EUR 1,000 (enough to cover the relevant household's average consumption) indicated, for instance, that – in addition to consumption expenses totaling EUR 900 – it had typically supported nonhousehold members with EUR 200 per month in the past year, the following message popped up:

"The sum of total consumption expenditure and regular remittances to nonhousehold members exceeds the household's total net income. Are the figures correct? If yes, please confirm the figure(s), or amend them as necessary."

The initial figures may in fact be confirmed in the cross-check, possible reasons being that the figures reported referred to different time periods, that the remittances were financed by the sale of assets, or that the household's income had since dropped as a result of one or more members losing their job. At any rate, inconsistencies would prompt the respondents to confirm or correct the total household income, remittances and consumption expenditure.

Other consistency checks programmed into the digital version of the HFCS questionnaire in Austria would allow the survey to proceed only once an answer identified as incorrect or inconsistent had been amended. However, these so-called "hard" checks were only used in cases where a particular answer could definitely be ruled out.

If individuals stated, for instance, that they had lived in Austria for 40 years but gave their age as 30, the following error message would appear:

"The respondent has been living in Austria for longer than his/her age allows. This is not possible. Please correct the information as necessary."

Thus, proceeding with the CAPI questionnaire required changing the age given to at least 40 years, or reducing the period of residence in Austria to 30 years or less (or changing both variables).

---

[4] A list of all the consistency checks that were programmed into the digital version of the questionnaire can be found in the online appendix.

4.4 Postinterview consistency checks

4.4.1 Expert data analysis

During the field phase of the fifth HFCS wave in Austria, household data deemed to be final by the survey company were forwarded to the OeNB in 24 batches. This means that the OeNB received household data roughly every two weeks during fieldwork. All batches of data were subjected promptly to expert data analysis.[5] On the one hand, these analyses served to improve the consistency of the data recorded for each household. On the other hand, they were used to check the survey software (in particular, to review the programming of the questionnaire) and the mechanisms used by the survey company to process the data.

The datasets for households actually interviewed and those for households that refused to participate were analyzed on a case-by-case basis. This made it possible to assess and optimize the success of interviewers in convincing households to participate. Hence, it was almost impossible for interviewers to cherry-pick "easy" or more readily accessible households, which would probably have created a bias toward certain households (e.g. those where housewives or pensioners live), thus distorting the data. The interviewers knew that the list of addresses was limited to the 7,636 households of the gross sample (see also chapter 6). This ensured that interviewers would not select the less difficult households and then move on to a new set of addresses. The incentive for interviewers to use the strictly limited address material as efficiently as possible was supported with a performance-related payment system and the relatively high effort that was required from interviewers to participate in the survey in the first place. Furthermore, area managers were advised to avoid allocating new households to interviewers before they had made sufficient effort to survey the households they were assigned at the time. The decision to exclude subsequent draws (substitute households) is among the key criteria for a successful survey, and is moreover essential for ensuring the representativeness of the sample (see e.g. Vehovar, 1999).

Initial analysis of the information on individual households during fieldwork covered the data provided on household structure, financial and real assets, debt and income, whether households had come into ownership of property by inheritance or gift, comments made by households or remarks made by interviewers, as well as the date, time and duration of the interviews. This information enabled a quick initial assessment of the interview's quality. The microdata on every single household were checked for consistency regarding their content and reviewed by at least two analysts from the OeNB HFCS team. Issues requiring clarification were discussed by the whole team, which then decided on the way forward.

In addition, this stage of the process was also used to assess the interviewers (see also chapter 3) and to address errors or misunderstandings. The shortcomings identified in this process were often minor in their nature, but interviewers whose results were not up to the required standards (e.g. regarding non-response) could be excluded.

4.4.2 Follow-up queries

If individual data analysis identified a problem but not how it could be corrected, households were contacted again by the survey company to clarify uncertainties and ensure that data were recorded correctly. Given the timely submission of interview results to the OeNB (around every two weeks) and the subsequent checks by the HFCS team, the survey company was able to address any queries to the surveyed

---

[5] The data were cross-checked against the results of the previous HFCS waves as well as external data sources such as the EU-SILC (conducted by Statistics Austria).

households promptly. A typical case of a data problem that was easy to spot and did not require queries was rewriting a negative sight account balance as a (positive) liability (overdrawn account) while setting the value of sight accounts to zero (see also section 4.6). This was simply a matter of adhering to the recording conventions as to where such liabilities should be recorded. Decisions on queries were always guided by the principle that any ex post data editing and the burden on participating households should be kept to a minimum. Many unusual results (e.g. particularly high asset values) were confirmed or corrected in the course of queries. All in all, follow-up queries (by phone) were necessary to confirm specific details of some 200 households. This is a percentage of households similar to that in the fourth wave.

4.4.3 Investigation of outliers

The checks on a case-by-case basis were aimed in particular at recognizing and processing outliers (exceptionally high or low values), which were recorded above all for wealth variables, the size of the household income or the size of the dwelling. Any outliers that were not removed from the dataset were generally not the result of interview errors but actually confirmed by the follow-up queries. Our recommendation for future studies based on HFCS data is therefore not to generally exclude outliers from the analysis, but rather to incorporate them in computations through the use of suitable methods.

4.4.4 Technical review of filtering and consistency

During the field phase, the consistency checks programmed into the digital version of the questionnaire and the rounds of expert data analysis were complemented with detailed automated consistency checks of the transmitted data.

All hard checks were applied repeatedly to the observations, for instance, in order to assess whether respondents might have given answers that precluded moving on to subsequent questions, thus requiring changes. The technical review also covered the questionnaire's complete set of filters to prevent programming errors leading to extensive and costly follow-up queries. Comprehensive tests of the questionnaire's programming prior to the start of fieldwork as well as a pilot survey of 50 households made it possible to largely exclude programming errors from the outset. Minor difficulties did occur, such as the double recording of verbatim information at one place in the questionnaire (see section 2.5.2.3). This deficiency was identified and corrected in a timely manner.[6] These filter checks also ensured that the coding of variables was consistent throughout the questionnaire.[7]

4.5 Flags

All edits (and imputations – see chapter 5) were documented with flag variables, which indicate how the individual HFCS observations were established (see table 3 for a list of the flags used to classify the observations). The flags used can be divided into five groups: flags used to identify recorded information (group I), incomplete or inadequate observations (group II), observations that were not recorded or later deleted for anonymization (group III), ex post edits (group IV) and imputed observations of different types (group V). To comply with requirements set out for the common (international) datasets covering all participating countries, some flags were aggregated for the international datasets (section 4.7).

---

[6] All deficiencies were corrected during the field period so that no household had to be excluded due to errors in the CAPI.

[7] All HFCS variables were assigned value labels that explain the coding. The coding of the individual variables is also included in the questionnaire (available in the online appendix).

Table 3

Flags used in the HFCS Austria

Group I	0	"Not applicable (i.e. skipped due to routing)"
	1	"Recorded as collected, complete observation"
	2	"Recorded as collected, but moved in iteration"
	12	"Recorded as found in other source, not collected in survey"
Group II	1050	"Not imputed, originally: Don't know"
	1051	"Not imputed, originally: No answer"
	1052	"Not imputed, originally not collected due to missing answer to a higher order question"
	1053	"Not imputed, originally collected from a range"
	1054	"Not imputed, collected value deleted"
	1055	"Not imputed, value not collected due to a CAPI error"
	1056	"Not imputed, set to missing due to incorrect answer to a higher-order question"
	1057	"Not imputed, collected value deleted but range information available"
	1058	"Not imputed, set to missing due to red button"
	1075	"Not imputed, specific answer code"
Group III	2050	"Missing, set to missing for anonymization purposes"
	2051	"Missing, set to missing because data were not collected"
Group IV	3050	"Edited, set to modified value as considered incorrect or unreliable"
	3051	"Edited, adjusted on the basis of other information obtained in the (national) survey"
	3052	"Edited, adjusted on the basis of verbatim records"
	3053	"Edited, set to missing (.)"
	3075	"Edited, set on the basis of follow-up with household"
	3076	"Edited, set on the basis of follow-up with interviewer"
Group V	4050	"Imputed, originally: Don't know"
	4051	"Imputed, originally: No answer"
	4052	"Imputed, originally not collected due to missing answer to a higher order question"
	4053	"Imputed, originally collected from a range"
	4054	"Imputed, collected value deleted"
	4055	"Imputed, value not collected due to a CAPI error"
	4056	"Imputed, originally value not recorded due to incorrect answer to a higher-order question"
	4057	"Imputed, collected value deleted but range information available"
	4058	"Imputed, set to missing due to red button"

Source: HFCS Austria 2023, OeNB.

Group I

The flags allocated to group I were used to identify recorded information. Specifically, all observations recorded during the interview were flagged "1" while all observations set to missing (".") were flagged "0." Information recorded in loops (see section 4.6.2.4) was paired with a flag of "2" if it had to be moved in the iteration of a loop. In other words, flag "2" observations were retained in the dataset exactly as they were recorded, but assigned a new iteration number. Observations flagged with "12" are not collected but taken from external data/information.

Group II

Recorded observations that were incomplete or inadequate were assigned group II flags. Such observations include cases where the respondent was unable or refused to answer the question (entries of "Don't know" or "No answer"), or proved unable to give a specific figure and provided a range instead. Included here are also observations that were not available on account of edits of either the variable in question or a head variable (flags "1054" and "1056"). If the edited observation was available as a range, it was assigned

a flag of "1057." If an observation was not available due to a CAPI error, it was given a flag of "1055." Observations that were not available because questions in a loop were skipped were flagged "1058" and special missing values were flagged "1075." In these cases, alternative information was collected.

For example, if gross income was unknown, but information on net income was provided, the variable for gross income was flagged "1075."

Observations with group II flags were not imputed (see chapter 5).

Group III

Group III flags identify observations and/or variables that were not recorded or that were recorded but later deleted from the datasets on account of anonymization requirements.

Group IV

Flags from group IV indicate an ex post edit of an entry. The following types of ex post edits can be distinguished: edits as a result of logical inconsistencies (flag "3050"); calculations that were adjusted using other information obtained in the survey, for instance with regard to life insurance contracts (see section 4.6.2.9 for details; flag "3051"); coding that was subsequently adjusted on the basis of verbatim records (see section 4.6.2.3; flag "3052"); edits made to delete a value and set the observation to missing, as in the case of duplicate entries (flag "3053"); and queries put to households (flag "3075") and interviewers (flag "3076").
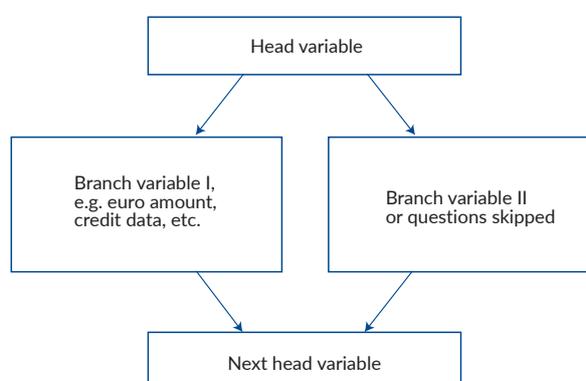
Group V

Flags from group V mirror those from group II. If it was possible to impute missing values, the first digit of the flag was changed to "4." For instance, if respondents had provided a range, which was subsequently imputed, rather than a specific figure, this observation was flagged with "4053" after multiple imputations. This ensures that all information can be tracked even after the imputations.

Figure 3 indicates how questions were typically structured in the HFCS questionnaire. Let us take employee income as an example to describe the structure of question blocks[8] and the use of flags.

For example, the head variable for recording employee income serves to ascertain whether or not a household has an income of this kind. If this yes/no question was answered with "Yes" the amount was recorded and the interview continued with the next head variable in the questionnaire – in this case, the question on self-employment income. If a household had no income of this kind, or if the respondent failed to provide the necessary information (i.e. responded

Figure 3

Sequence of questions

Source: HFCS Austria 2023, OeNB.

---

[8] See chapter 2 for details of the structure of the whole questionnaire.

with "Don't know" or "No answer"), the interview continued with the question on self-employment income (the next head variable). Depending on which answers were given, all the observations recorded were initially flagged "1" or "0." If the response to a subsequent question (e.g. on employment) revealed that a "No" given the question on employee income was in fact incorrect, the initial response was corrected and flagged "3050" (edited, set to modified value as considered incorrect or unreliable) and the corresponding data entry field for the value was flagged for imputation. Following imputation, the value was then reflagged "4056" (imputed, original value not recorded due to incorrect answer to a higher-order question).

Or, if the question on a household member's highest education qualification (variable (A)PA0200) was answered by selecting the category "Other qualification" and if that answer was subsequently found to match one of the predefined categories, the observation was flagged "3052" (edited, adjusted on the basis of the verbatim records) in the flag variable of the individual dataset.

This flag system allows the origin of observations in the HFCS to be tracked. In the Austrian questionnaire the section on pension policies was structured differently to reflect national circumstances and to be consistent with previous waves. This section was adapted ex post to fit the requirements for the international data.

To allow for the merging of datasets, no flags were used to encode the variables for identifying households and individuals, nor were the country codes and the imputation's iteration number flagged. The flags described here provide for a more detailed breakdown by category than those incorporated into the international HFCS dataset, which can be obtained from the ECB. For reasons of international consistency, the flags were aggregated prior to being submitted to the ECB (see section 4.7).

4.6 Ex post editing

4.6.1 Case-by-case review

A detailed case-by-case review of all households allowed inconsistencies to be identified and eliminated through follow-up queries and ex post editing. Specifically, respondents' answers were checked for plausibility against known benchmarks, including descriptive statistics (e.g. on average income) compiled on the basis of completed HFCS interviews and external sources of data. Moreover, the review process heavily relied on auxiliary variables that recorded information in aggregated form and/or in a variety of other ways.

Both interviewers that produced nonstandard results (see chapter 3) and follow-up queries made by the survey company were reviewed in particular detail.

Expert assessments were generally used to resolve the following issues through ex post edits:

- Double entries: Cases where an inheritance, for instance, was recorded under both "Household main residence inherited" and in the "Inheritances received" chapter, or where the same income was recorded in two different income categories, had to be corrected.
- Missing or additional "zeros": In a few cases interviewers added or left out a zero by accident when recording amounts; this had to be amended accordingly.
- Implausible values: Values that remained implausible after follow-up queries had to be set to missing and were subsequently imputed.
- Often, information could be gained from the many additional comments made by respondents. If the additional comments made it necessary to change the collected information, changes were made respectively.

- Data entry errors by interviewers: For instance, a wrongly recorded date of a contact attempt by an interviewer was easy to correct because of the potential other contact attempts and the immediate data transmission of completed households.
- Also, all data obtained through follow-up queries were used in this step to correct individual observations in the dataset where necessary.

Such edits related to the whole questionnaire, not just to individual variables. Amendments to recorded data were kept to a minimum and – wherever follow-up queries and/or the use of auxiliary variables (such as verbatim records) failed to provide further information – inconsistent observations were set to missing and flagged for imputation. Inconsistent or implausible observations were processed with great care and only deleted if there was absolutely no doubt about the inconsistency.

4.6.2 Structural editing

4.6.2.1 Data cleanup

When answering the HFCS questions, respondents occasionally gave inaccurate answers but subsequently corrected those answers as they proceeded through the questionnaire. These corrections also necessitated a change in the sequence of questions following the initial question because the new answers called for different filter settings. The "wrong" initial path through the questionnaire, however, remained in place for transparency reasons and had to be subsequently corrected.

4.6.2.2 Currency conversion

Respondents could specify any amount in various currencies (see chapter 2). The edits set out below relate both to specific amounts and ranges indicated by the respondents (predefined ranges had to be specified in euro).

Typically, amounts were given either in euro or in Austrian schillings. In particular, the value of the main residence (both the purchase price and the current value) was often given in Austrian schillings. All Austrian schilling amounts were subsequently converted into euro at the irrevocably fixed conversion rate of EUR 1 = ATS 13.7603.[9] Some amounts were also given in Deutsche mark (DEM). These amounts were also converted at the irrevocably fixed conversion rate, namely EUR 1 = DEM 1.95583.

In a few cases amounts were also given in Swiss francs and pound sterling. For the amounts outstanding on a loan and the current value of a property the average exchange rate for the month in which the interview took place were used.[10] For income in a foreign currency, the average exchange rate for the year 2023 was applied.[11]

There were cases in which the given value originated in a time where the euro had not yet been introduced and this value was given in a currency other than Austrian schillings. These cases concerned initial values of real estate and values of inheritances. Here, the amount was first converted into Austrian schillings on the basis of the exchange rate applicable at the time and then from Austrian schillings into euro according to the fixed ATS/EUR exchange rate.

---

[9] See www.oenb.at/isaweb/report.do;jsessionid=31767F3B9E6FA661A8A4CD5CB700B5A7?report=2.12 (accessed on July 21, 2025).

[10] See https://www.oenb.at/Statistik/Standardisierte-Tabellen/zinssaetze-und-wechselkurse/Wechselkurse.html (accessed on July 21, 2025).

[11] See http://fxtop.com/en/historical-exchange-rates.php?MA=1 (accessed on July 21, 2025) for historical exchange rates.

4.6.2.3 Verbatim records

For many questions, respondents were given the option of choosing the category "Other" and providing a verbatim response, mainly with a view to making the questionnaire as user-friendly as possible. Thus, a verbatim description could be recorded if it was not possible to assign a respondent's answer to a predefined category during the interview. The verbatim entries were used to assign answers to specific categories ex post, which proved to be possible in the majority of cases. Wherever this could not be done, the initial categorization of the observation as "Other" was retained. Some data, such as data on the occupation (ISCO coding in the variable PE0300) of an employed individual or the main activity (NACE coding in the variable PE0400) of the company where the individual is employed, were collected entirely in verbatim form and coded ex post. All observations subjected to ex post edits on the basis of verbatim records were flagged "3052" (see section 4.5 for details on the flags).

4.6.2.4 Navigation of loops

As outlined in detail in section 2.6.1, some pieces of information were recorded in loops, which required interviewers to run through an identical set of questions for each individual item from a group of items owned by the household. Information on the following items was collected using loops:

- mortgages on the main residence
- real estate assets apart from the main residence
- mortgages secured against these other properties
- unsecured loans from family and friends
- other unsecured loans
- businesses owned by the household
- life insurance policies
- inheritances and gifts

Below we provide an explanation of the edits which were required because of loop questioning.

Recording sequence

The sequence of items that were covered in loops followed a predefined order. With regard to mortgages secured against the main residence, for instance, the first iteration of questions related to the mortgage with the highest amount outstanding, the second iteration of the loop to the mortgage with the second-highest outstanding amount and the third iteration to the third-highest loan amount outstanding. Some respondents did not always adhere to this sequence. Such cases were recoded in the editing process – with the exception of the loop questions on inheritances, for which no recoding was carried out because respondents were prompted to discuss the inheritances received in descending order of relevance for the household's current wealth situation. Moreover, they were instructed to indicate amounts as transferred rather than current amounts. After all, certain inheritances could have gained (or lost) more in value than others since the inheritance date; or inherited residential property might since have been passed on to children, causing it to be irrelevant for the household's wealth situation at the time of the interview.

Every variable within a loop that was replaced with observations recorded for the same variable in another iteration was flagged "2" (see section 4.5). Wherever a variable set to missing in one iteration was replaced with the same variable set to missing in another iteration, it was flagged "0" (Not applicable (filtered out)).

Skipping questions

In order to avoid breaking off an interview in mid-loop, respondents were allowed to skip parts of loop questions and to proceed directly to the summary questions, where either the residual sum total of the not yet recorded loans and/or businesses (more than three loans or businesses) or the sum total of all loans and/or businesses was recorded. If questions within the loop for inheritances and gifts were skipped, information on the sum total of all inheritances was always requested in the summary question. As the summary questions from all sections of the dataset to be sent to the ECB were supposed to cover only any items that went beyond the first three itemized loans, real estate assets besides the household main residence and private businesses, the relevant summary responses had to be edited accordingly. For ease of reference, examples of such edits are described below on the basis of the section of the questionnaire dealing with other unsecured loans (see section 2.5).

In the three cases in which a household had taken out only one unsecured loan and had skipped questions within a loop, the type of edit depended on whether the respondent had (1) indicated the outstanding amount only in response to the summary question; or (2) both when going through the first loop of questions and in answering the summary question; or (3) neither during the first loop of questions nor in response to the summary question. If the respondent had indicated the outstanding amount only in response to the summary question (variant 1), this amount was entered as the answer to the appropriate question (in the first loop) and the entry under the summary question was set to missing. If the respondent had indicated identical amounts in answering both the loop and the summary question (variant 2), the latter was set to missing since it was a duplicate entry.[12] Where no amount was given at all, neither within the loop nor in the summary (variant 3), only the summary question was set to missing.

In cases where a household had taken out two unsecured loans and had skipped questions within a loop,[13] the type of edit depended on whether the respondent had (1) specified the value of the highest outstanding loan and indicated an aggregate amount in response to the summary question; or (2) indicated outstanding amounts in response to both question loops and the summary question; or (3) specified an amount only in the answer to the summary question; or (4) given no amounts at all, neither in the answers to the loop item questions nor in the answer to the summary question.

If variant 1 was the case, the amount outstanding for the lower of the two loans was taken to be the difference between the amount given in the answer to the summary question and that given in the first loop. This, however, was only done if the sum total of the two outstanding loans exceeded the amount outstanding from the first loan. If it was lower, it was assumed that the amount given in the answer to the summary question was not the sum total of the two outstanding loans, but rather the amount outstanding for the second loan. In both instances, the summary question was subsequently set to missing. If variant 2 was the case, the amount given in response to the summary question was set to missing. If only the sum total of the two outstanding loans was given (variant 3), it was used as the upper bound for both the first and the second loan for the imputation model. This was the case for one household that held two other unsecured loans and had skipped some loop questions.

---

[12] Where the amounts given were not identical, the one specified in response to the loop questions on the first loan was deemed to be more relevant than that given as an answer to the summary question. The reasoning behind this procedure is that the loop questions relating to the first loan contained a question explicitly asking for the amount outstanding on an unsecured loan, so the amount given there was regarded as more trustworthy.

[13] In the fifth wave of the HFCS in Austria, no household indicated to have two outstanding loans of this type and skipped the questions within the loop. The explanation of potential editing, however, is kept due to the cases with more than two such loans and their similar treatment.

If no amounts were given at all, neither in response to the loop questions for each of the two outstanding loans, nor in answer to the summary question (variant 4), the summary question was set to missing.

The editing procedure followed in cases with three loans and skipped loop questions prior to the recording of the individual amounts outstanding was similar to that used for two loans when loop questions were skipped.[14]

All edits were again flagged correspondingly. The example laid out above covers all potential variants. Not all these variants are necessarily seen in every wave of the HFCS in Austria.

Summary questions

Every loop of questions ended with summary questions (see figure 2 in chapter 2). The variables for these questions exclusively contained information on any additional items above three (in some instances five iterations) per household. As indicated in figure 2, the summary questions were ultimately also put to all respondents who had refused to indicate the number of a given item in the household. In such cases of nonresponse, the information provided here was used for multiple imputations (chapter 5) and deleted from the dataset ex post.

4.6.2.5 Sight account balances and overdrafts
A few households misreported a negative balance on their household sight account as a negative value of sight accounts (HD1110). For this, however, a separate variable was available. In this area, there were also occasional duplicate entries as well as misplaced entries that subsequently had to be edited.

4.6.2.6 Rent variables
The HFCS questionnaire included questions on the amount of housing rent paid both excluding and including utilities. In the case of some households, the given rent excluding utilities was higher than, or equal to, rent including such costs, which is logically impossible as housing cannot be "run" free of charge. Some of these households entered just the utility costs under the item "Rent including utilities." In the course of editing, these were added to the amount entered under "Rent excluding utilities" to obtain the "Rent including utilities." In the case of other households, the "Rent including utilities" was set to missing and flagged for imputation, with the "Rent excluding utilities" serving as the lower bound to the "Rent including utilities."

In addition, the item "Rent including utilities" was set as the upper bound for the variable "Rent excluding utilities" and used for imputations whenever the answer to the latter was not an amount (i.e. read "Don't know," "No answer" or "Rent excluding utilities unknown") (see also section 5.4.6 on the use of bounds in the imputations).

4.6.2.7 Agricultural businesses
As defined in the HFCS, farmers are owners of an agricultural business. Separating the asset components of households that own an agricultural business sometimes posed a problem to respondents, in particular with regard to their main residence and the investments in their business. Such cases, therefore, had to be analyzed separately. In this context, the extra questions and guidance added to the questionnaire for the second wave and kept for the following waves (see also section 2.6.3) proved very helpful during the various steps of data processing.

---

[14] In the fifth wave in Austria there were two such cases.

Very few farmers did not report their agricultural business as an investment in self-employment businesses. For these households, data on investments in a self-employment business had to be imputed. The NACE code for such businesses was set to that for "agricultural businesses," and at least the individual who stated that he/she worked as a farmer was deemed to be employed in this agricultural business. The legal form of the respective business was edited to read "sole proprietorship." The additional guidance implemented for the second wave was again used in the fourth and continued to keep the number of such cases low.

For all farmers, additional auxiliary variables were created for the combined value of the main residence and the agricultural business (business assets) as well as for the main residence's share in this amount. For households that were not able to separate their assets and specify the share themselves, information on the total value and on the main residence's share was used. For households that had specified both the value of their main residence and that of their private business, as required, the combined value and the share of the main residence was calculated. If information was partially missing, it was flagged for imputation (see section 5.3).

The category of agricultural businesses was subjected to case-by-case reviews. Particularly complex cases were clarified through follow-up queries and corrected where necessary.

4.6.2.8 Individual variables for investments in self-employment businesses

The variables for household members employed in a business owned by the household were edited as follows:

To be able to cover even unusually large households, variables were created for up to 18 individuals per household for the CAPI version of the questionnaire. The largest household successfully interviewed in Austria had only 10 members, however, so all variables in excess of that number were deleted from the dataset. Moreover, the coding was changed from yes/no questions for each household member (the type of coding used in Austria) to a single variable in the international user database containing only the number of persons within a household working in the particular business.

At the same time, all NACE codes for household members employed in the business were checked against the information contained in the P-file and corrected where necessary.

4.6.2.9 Life insurance policies

Information on assets held in life insurance contracts was recorded through questions ensuring that the answers were both as precise as possible and prone to only a few errors. There was, for instance, no direct question on the total value of such assets, but rather a series of questions on the start of payments, the frequency of payments (monthly, yearly or single payment), the type of life insurance (benefits to be provided at the death of the policy holder or at a given date, or a hybrid form) and the amount of the current payments for every single life insurance contract in the household. For all life insurance policies with a set payout date and/or all hybrid policies, the value of the assets held in life insurance contracts was calculated as the sum total of all payments. In cases where one or several details were not given, the remaining observations were used as bounds for the values to be imputed. Insurance policies which do not pay out capital if the insured lives beyond the term period do not constitute wealth; they were therefore excluded from this calculation.

4.6.2.10 Income variables

The following categories (variable name in parentheses) of personal income were recorded separately for every member of the household who was 16 years of age or older:

- employee income (PG0100 and PG0110)
- income from self-employment (PG0200 and PG0210)
- income from public pensions (PG0300 and PG0310)
- income from private and occupational pension plans (PG0400 and PG0410)
- income from unemployment benefits (PG0500 and PG0510)

This information was supplemented by the following income categories that were recorded per household:

- income from public social transfers (HG0100 and HG0110)
- income from private transfers (HG0200 and HG0210)
- income from other private transfers (HG0250 and HG0260)
- income from real estate assets (HG0300 and HG0310)
- income from financial investments (HG0400 and HG0410)
- income from investments in self-employment businesses or partnerships (HG0500 and HG0510)
- income from other sources (HG0600 and HG0610)

In the case of the first four personal income categories, respondents could indicate their net income if they did not recall their gross annual income (see chapter 2). Likewise, respondents could indicate their net income from financial investments if they did not know their gross income in this category.

Where only a net amount was entered for individual incomes, the gross income was calculated with the aid of the Austrian finance ministry's gross-to-net calculator,[15] based on information on the type of income, the structure of the household (with reference to the tax credits for single parents and single earners), the employment status and age of any children, the province and the respondent's employment status (employed (holding a "blue collar" manual job; or an office, sales or services job) or retired).[16] Wherever both parents were gainfully employed, the single earner's tax credit was assigned to the main earner, i.e. the parent with the higher income (as long as the legal requirements were fulfilled and the partner did not earn more than EUR 6,000 per annum).

Given the far greater scope for tax deductions for self-employed people, the gross-to-net conversion of income from self-employment was not generally based on the precise figures. Precise conversions were made only for gross annual incomes below EUR 11,000, which are tax-free, albeit sometimes subject to social insurance costs. For such cases the gross amount was set equal to the net. For all other gross (or net) incomes from self-employment (36 individuals),[17] a range was created for the purpose of imputing specific amounts by adding EUR 10,000 to and by subtracting EUR 10,000 from the gross (or net) amount converted subject to the conditions for employees in office, sales or service jobs. This range reflected the uncertainty that such a conversion entails, without losing the important information of the actual range within which the value is placed. This range then was used as a bound in the imputation of a precise value (see section 5.4.3).

---

[15] See https://onlinerechner.haude.at/BMF-Brutto-Netto-Rechner/ (accessed on July 21, 2025). Available in German only.

[16] Apprentices were categorized as manual workers in the conversion, while civil servants were treated as employees in office, sales or service jobs on grounds of their more favorable taxation.

[17] These cases include individuals whose income is sourced at least partly from self-employment. Cases where there are further sources of income have to be treated accordingly as is discussed later in this section.

To calculate the gross income from financial investments, 25% withholding tax was added to amounts given for net income.

When individuals reported more than one income from more than one type of employment, a slightly more complex method for calculating the gross and net income values had to be applied to account for the complexities of the Austrian tax system for such situations. For instance, individuals could receive their main income from a job with the status of employee and additionally earn some money from self-employed work. Other common cases included pensioners who topped up their public pension with income emanating from self-employment or also employment. When converting gross income to net, all forms of income of an individual are taken into account. Sometimes respondents reported their net income from one source and their gross income for another. The gross and net values were approximated by converting one of the values to its counterpart gross or net value, applying the conversion to the sum of both net or both gross and using this amount to calculate the still missing net or gross value of the income not used in the initial step.

If the net amount was only recorded as a range, the upper and lower bounds were converted into gross values that were subsequently used in the imputations. All converted values were flagged "3051."

Using flags as a basis, table 4 gives an indication of the number of edits relating to employee income. The table also illustrates the use of flag variables (see also section 4.5).

The question on the amount of employee income received (variable PG0110) was put to a total of 2,267 individuals. 976 respondents (43.1%) expressed their annual income in gross terms. A further 50 respondents (2.2%) answered "Don't know" and 44 individuals (1.9%) opted for "No answer." 175 respondents (7.7%) specified their income amount using a range. For 53 individuals, a range could be calculated from other information given, i.e. the range was given for net income and converted. The responses of 8 individuals (about 0.5%) were edited and if necessary set to missing and flagged for imputation. 899 of the respondents (39.7%) provided their net income, which was then converted as described earlier. The responses of the remaining 8 individuals (around 0.4%) were corrected on the basis of follow-up queries.

Table 4

Number and share of edits of gross employee income based on flags

	Number of persons	Share
		%
Number of persons receiving employee income	2,267	100
Answer recorded, complete observation (flag 1)	976	43.1
Not imputed, originally: „Don't know" (flag 1050)	50	2.2
Not imputed, originally: „No answer" (flag 1051)	44	1.9
Not imputed, originally not collected due to higher order missing (flag 1052)	7	0.3
Not imputed, originally collected from a range (flag 1053)	175	7.7
Not imputed, collected value deleted (flag 1054)	5	0.2
Not imputed, set to missing due to incorrect answer to higher order question (flag1056)	8	0.4
Not imputed, collected value deleted, but range information available (flag 1057)	53	2.3
Edited, set to modified value as considered incorrect or unrealiable (flag 3050)	41	1.8
Edited, adjusted on the basis of other information obtained in the (national) survey (flag 3051)	899	39.7
Edited, adjusted on the basis of verbatim records (flag 3052)	1	0.0
Edited, set on the basis of follow-up with the household (flag 3075)	8	0.4

Source: HFCS Austria 2023, OeNB.

4.6.2.11 ISCO and NACE classification

As required by the euro area blueprint questionnaire, the main occupation of respondents was recorded (in variable PE0300) using the occupation codes and titles set out in the International Standard Classification of Occupations (ISCO-08). Making individual members of each household classify their jobs themselves, however, would have been extremely difficult for respondents without any advance knowledge of the ISCO codes, possibly giving rise to misclassifications. Therefore, verbatim answers were recorded for the Austrian HFCS question on job titles and/or main job tasks. That information was later paired with the corresponding ISCO-08 codes, as published by Statistics Austria (in German).[18] As required by the ECB, classification was based on the two-digit ISCO codes (major subgroups). To this end, the verbatim record of the job title and related main tasks was supplemented with individual data relevant for ISCO classification (in particular, the respondent's qualifications and the main activity of the company where the respondent worked). The variable PE0300 to be submitted to the ECB was first flagged "3051" (edited, adjusted on the basis of other information obtained in the (national) survey) and aggregated in a next step (see section 4.7).

Also, the main activity of the company (PE0400) where the respondent worked was first recorded verbatim and then assigned a single-digit NACE rev. 2 code.[19]

4.6.2.12 Highest education qualification

To account for developments in the International Standard Classification of Education (ISCED), the highest education qualification of all household members has been recorded in substantially more detail since the second HFCS wave. Respondents were not asked about ISCED categories, however, but were prompted to indicate their qualification based on Austria's academic degree hierarchy (bachelor, master (including previous versions) and doctorate).

4.6.2.13 Exclusion of successful interviews

The final data do not include 20 households with a successful interview which were dropped from the data because they either did not belong to the target population or the quality in terms of item non-response did not meet the required standard.

4.6.2.14 CAPI errors encountered with the questionnaires

Although very few flaws in the CAPI remained in the beginning of the field period, they did not necessitate any edits. This shows, in particular, improvements over previous waves.

4.7 Formatting and editing after multiple imputations

Any information collected at a greater degree of granularity in Austria than required for the international dataset was processed further upon imputation so as to bring the level of aggregation in line with the international requirements. The most important aggregations can be summarized as follows:

- Marital status: The categories "Married and living together with spouse" and "Married, but separated" were aggregated as "Married."

---

[18] For further information, see the Statistics Austria classification database https://www.statistik.at/en/databases/classification-database (accessed on August 5, 2025) in the "Occupations" section.

[19] For further information, see the Statistics Austria classification database https://www.statistik.at/en/databases/classification-database (accessed on August 5, 2025) in the "Economic activities" section.

- Education: Categories specific to Austria were paired with ISCED[20] 2011 codes[21] and classified as ISCED level 0 ("Early childhood education or no education")[22]; ISCED level 1 ("Primary education"); ISCED level 2 ("Lower secondary education"); ISCED level 3 ("Upper secondary education"); ISCED level 4 ("Post-secondary nontertiary education"); ISCED level 5 ("Short-cycle tertiary education"); ISCED level 6 ("Bachelor's or equivalent level"), ISCED level 7 ("Master's or equivalent level") and finally ISCED level 8 ("Doctoral or equivalent level").
- Employment status/relationship: More detailed categories were aggregated.
- Main residence – tenure status: More detailed categories were aggregated.
- Reasons for refinancing: With regard to collateralized loans, "For the conversion of a foreign currency loan" was available in Austria as an additional category for this variable. This category was added to "Other" in the international dataset.
- Use of additional real estate property: In this variable, "Buy-to-let apartment" was available as an additional category in Austria; in the international core dataset, this category was added to "Other."
- For the international version the question on loans secured with further property was structured differently than in the previous waves in that loans secured by real estate were asked together with different characteristics of the property instead of in a separate loop as before. The Austrian questionnaire kept the previous structure of asking about further properties and any loans secured with these in two separate loops. To bring the data into the structure required by the international dataset, a question on which property was used to secure a particular loan was used to link the two loops.
- Due to the high complexity of bringing the flags of the national data into the structure required by the international dataset, the flags in the international dataset corresponding to the variables concerning loans secured by other real estate next to the household's main residence were simplified by using the flag value "0" if the national value was missing and "13" otherwise.
- Purpose of private and noncollateralized loans: The categories "To finance a deposit for the housing association" and "To support friends and family" were allocated to "Other."
- Rejection of a loan application: Information of this type was recorded in three variables with multiple responses; it was then aggregated into two variables.
- Business' legal form: More detailed categories were aggregated.
- Money in savings plans with building and loan associations and life insurance policies: Data recorded on these two investment methods are aggregated into savings (HD1200 and HD1210).
- National information regarding crypto assets (variables ahd1720-ahd1770) are aggregated to other financial assets (hd1900 to hd1920) in the international core user database.
- Type of assets received (survey questions on inheritances and gifts, HH030$a-i): The sequence based on values was abandoned.
- Provider of assets (survey questions on inheritances and gifts): More detailed categories were aggregated.
- Five variables of the questions concerning impatience ((a)hiz050x) were aggregated – without loss of additional information – to adhere to international standards.

---

[20] International Standard Classification of Education.

[21] For explanations, see https://www.klassifikationsdatenbank.at/KDBWeb/kdb.do?FAM=BILD&&NAV=EN&&KDBtoken=null (accessed on July 21, 2025).

[22] Only very few of the participants (11 individuals) in the fifth wave in Austria had completed less than compulsory education.

- Purpose of saving: The sequence ordered by relevance was abandoned.
- Type of dwelling (paradata): More detailed categories were aggregated.

In Austria, more specific flags were used in some areas (see also section 4.5). To conform to international standards, these flags were, in general, aggregated as follows:

- flag "1051" was recoded as "1050"
- flags "1052", "1055", "1056" and "1057" were recoded as "0"
- flag "1054" was recoded as "1050"
- flag "1057" was recoded as "0"
- flag "1058" was recoded as "1050"
- flags "3075", "3076", "1075" and 2 were recoded as "1"
- flags "4055" and "4056" were recoded as "4052"
- flag "3051" was recoded as "1"
- flag "3052" was recoded as "1"
- flag "3053" was recoded as "0"
- flag "4051" was recoded as "4050"
- flag "4057" was recoded as "4052"
- flag "4058" was recoded as "4050"

The following variable is an exception to this rule:

- The variable featuring rent excluding utilities (HB2300) was flagged "1075" to identify the special cases when only the rent including utilities was known to respondents. These costs were subsequently imputed, with the rent including utilities serving as bounds, and reflagged "4053."

In analogy to the first waves, some of the additional data over and above those of the ECB's HFCS datasets, which are collected at the national level and contain all the variables specified by the ECB, will probably be available from the OeNB as of winter 2026. The additional information includes additional variables, as well as a more detailed breakdown of certain variables. Datasets may be merged on the basis of both identification numbers and imputation numbers. The aggregation of flags in Austria yields less information for Austrian data in the internationally available data. For more information regarding flags, please contact the HFCS team of the OeNB.

4.8 Concluding remarks and online appendix

The underlying rationale of editing was to edit only those observations that had clearly not been recorded correctly. In cases of ambiguity, the first possibility considered was to conduct ex post checks by phone. This option allowed many observations either to be corrected or to be confirmed.

Knowledge of the steps undertaken to check the consistency of the data is essential both for analyzing the data and for understanding how the observations originated. In addition, the use of flags makes it possible for users to develop an imputation model of their own, to dispense with imputations, or to resolve the problem of item nonresponse in another way.

The online appendix, which contains the information provided here on the edits and consistency checks applied in the HFCS in Austria, includes a list of the consistency checks programmed into the digital version of the questionnaire.[23]

---

[23] All documents included in the online appendix are available at https://www.hfcs.at/en/results-tables/.

5 Multiple imputations

5.1 Introduction

A common problem with voluntary surveys is item nonresponse, i.e. the fact that some survey participants do not answer all questions.[1] This is especially the case with surveys that pose complicated or sensitive questions (e.g. about income or wealth).

Disregarding the problem of missing information due to item nonresponse would lead to biased estimates. For the HFCS data, we therefore used multiple imputation with chained equations.

The idea behind this approach is to substitute missing values in the dataset with several values that have been estimated based on an iterative Bayesian model. The main aim of this procedure is to impute in such a way that the associations between all variables are preserved in terms of maintaining the correlation structure of the dataset. Under this approach, the missing values of each variable are estimated by taking into account a maximum number of available variables. To account for the uncertainty of the missing values, not just one value per missing value is imputed, but several (in the case of the HFCS, five).

This data imputation approach is also used by similar surveys, such as the U.S. Survey of Consumer Finances (SCF – see Kennickell, 1998) and the Spanish Survey of Household Finances (EFF – see Barceló, 2006).

As multiple imputation is a very time-consuming process, most institutions that carry out surveys, including the HFCS, provide users with datasets which are already imputed. This ensures that all users can work with the same imputed datasets. In the case of the HFCS, users can identify every imputed value of any variable by looking at the corresponding flag variable (section 4.5). Thus, they have the possibility to carry out nonresponse analyses or imputations on their own, or to use other methods for dealing with item nonresponse in their analyses.

This chapter is structured as follows: In section 5.2, we present data on item nonresponse in the HFCS. Section 5.3 describes the imputation procedure used, and in section 5.4 we explain the imputation model specification and how the imputations were executed. Finally, some imputation results are presented in section 5.5.

5.2 Item nonresponse

Table 5 shows selected statistics on item nonresponse for the interviews from the fifth wave. On average, each household has 16.7 missing values, which means that item nonresponse was limited to 0.8% of all the questions (variables) addressed to each household. However, the respective percentage for the euro variables amounts to 3.7%. This suggests that questions of this kind might be perceived as sensitive or difficult to answer. These figures are comparable to the ones of the last wave but still significantly below the level of the previous waves (see the methodological report of the last wave for more information on the changes due to COVID-19).

There are different ways of analyzing datasets that include variables with missing values.[2] In most statistical packages, the default method is the complete-case analysis method. This method entails deleting all households that have missing values in any of the variables of interest and basing the analyses solely

---

[1] A common related problem that occurs in surveys is unit nonresponse, which means that no questions are answered at all because, for example, a household declined to take part in the survey. This problem is addressed with the construction of HFCS nonresponse weights (chapter 7).

[2] For a comprehensive study, see Little and Rubin (2019).

Table 5

Item nonresponse per household in wave 5 (unweighted)

	Mean	Median	Minimum	Maximum
Number of variables asked				
All variables	1,964.6	2,010.0	1,337	2,222
Euro variables	123.0	129.0	36	156
Number of variables with missing values				
All variables	16.7	10.0	0	425
Euro variables	4.5	3.0	0	51
Share of variables with missing values in %				
All variables	0.8	0.5	0	19
Euro variables	3.7	2.1	0	53.4

Source: HFCS Austria 2023, OeNB.

Note: Interval responses are considered as missing values with regard to the corresponding euro variable and are not included as a separate variable. A question addressed to several household members is entered as several variables, one for each household member.

on complete observations. However, the loss of information resulting from this method leads to two problems: First, it biases estimates if complete observations differ systematically from incomplete ones; second, even if an estimate is unbiased, the estimation would be less precise due to the observations lost. To illustrate how significant the loss of information would be in the case of the HFCS, table 6 shows item nonresponse rates across some selected variables.

For example, when asked about the value of their main residence, 76.1% of households provided a specific amount (column 3 of table 6). The other 23.9% of households are item nonrespondents: Either they provided a (prespecified or individual) range (19.7%, column 4), responded with "Don't know" or "No answer" (4.1%, column 5) or their response was set to missing[3] (0.2%, column 6) in the editing process. Nonresponse rates[4] vary substantially across items. Variables with high nonresponse rates include e.g. questions related to the value of publicly traded shares (100% – 71.7% = 28.3%) and the household's gross income from financial investments (100% – 45% = 55%). However, 17.9% and 40.3% of households provided at least a range for these two questions, which confirms the importance of asking range questions when a euro question remained unanswered. Range questions provide valuable and often very precise information (see the online appendix and section 2.6.2 for the questionnaire and information on the design of euro loops). Variables with low nonresponse rates include variables such as employment status (100% – 100% = 0%) and the monthly amount paid as rent (100% – 97.4% = 2.6%).

Table 6 (column 2) also shows another aspect of item nonresponse in the HFCS: There are variables known as branch variables (see chart 3 in chapter 4) that may also have missing values due to nonresponses to a higher-order question (head variable) and that are thus set to missing. For example, before the euro question on gross income from financial investments is asked, households are asked a yes/no question determining whether they have this type of income or not. Only those that answer affirmatively (59.5%) are then asked the question on the amount of income; the other households, including the 7.9% of households that did not answer the yes/no question, automatically skip the euro question. As it is unknown, however, whether the 7.9% of households that did not answer the yes/no question have a positive gross income from financial investments or not, their nonresponses must also be considered as second-order (or higher-order) missing values when analyzing nonresponse to a euro question.

[3] See chapter 4 for more details.

[4] The nonresponse rate is calculated by subtracting the value in the "amount" column (column 3) in table 6 from 100%.

Table 6

Item nonresponse for selected variables in wave 5 (unweighted)

	Household has item		Responses by households that have the item			
	Yes	Unknown	Amount	Range	"Don't know"/ "No answer"	Other missing values[1]
	(1)	(2)	(3)	(4)	(5)	(6)
	%					
Value of main residence[2]	38.9	0.0	76.1	19.7	4.1	0.2
HMR mortgage 1: amount still owed	10.7	0.2	81.6	8.9	9.5	0.0
Monthly amount paid as rent	54.1	0.0	97.4	2.4	0.2	0.0
Other property 1: current value	10.0	0.1	71.8	15.5	8.8	3.9
Other property mortgage 1: amount still owed	1.2	0.0	87.9	6.1	6.1	0.0
Value of sight accounts	99.5	0.2	78.2	10.1	11.4	0.3
Value of saving accounts	98.4	1.6	77.2	10.8	8.5	3.4
Value of publicly traded shares	5.1	0.1	71.7	17.9	10.3	0.0
Amount owed to household	5.6	0.1	90.6	8.1	1.3	0.0
Employment status (main activity) (person 1)	100.0	0.0	100.0	0.0	0.0	0.0
Gross employee income (person 1)	46.3	0.0	89.9	7.7	2.0	0.4
Gross income from unemployment benefits (person 1)	4.2	0.0	88.4	8.3	3.3	0.0
Gross income from financial investments	59.5	7.9	45.0	40.3	14.3	0.4
Gift/inheritance 1: value	36.2	0.7	76.1	10.4	12.2	1.4
Amount spent on food at home	100.0	0.0	93.2	6.8	0.0	0.0

Source: HFCS Austria 2023, OeNB.

[1] Missing values due to editing measures and exits from loops.
[2] Based on the HB0900 variable.

Note: HMR = household main residence.

Table 7

Thus, if a complete-case analysis were to be carried out with the HFCS data, the loss of information and the resulting loss in precision of unbiased estimates would be considerable owing to the large amount of variables with higher-order missing values. Furthermore, as complete observations usually differ systematically from incomplete ones, complete-case analysis would bias the estimates.

For illustrative purposes, table 7 shows a regression of nonresponse to the question regarding the balance of sight accounts ("1" if the value is missing, "0" otherwise) for several explanatory variables. We can see that item respondents differ significantly from item nonrespondents, because respondents tend to be male, tend to live in smaller main residences and smaller households and tend not to live in Vienna. Thus, a complete-case analysis of the value of sight accounts would bias the estimates toward a population with these household characteristics.

Logit regression of nonresponse in the euro question on value of sight accounts (unweighted)

Covariates	Coefficient
Female (person 1)	0.186*
	(0.0954)
Age (person 1)	−0.00190
	(0.00373)
Tertiary education level (person 1)	0.105
	(0.117)
Employed/self-employed (person 1)	−0.00848
	(0.129)
Residence is in Vienna	0.317***
	(0.112)
Size of main residence	0.00226**
	(0.000952)
Household size	0.0823
	(0.0502)
Constant	−1.769***
	(0.287)
Observations[1]	2,816

Source: HFCS Austria 2023, OeNB.

[1] The remaining 33 observations of the dataset show missing values in one of the covariates and/or filter missing remarks in the dependent variable and are thus not included in the regression.

Note: Standard errors in parentheses; *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

5.3 HFCS imputation procedure

To impute HFCS data, we have chosen a procedure implemented in the statistics software Stata by Royston (2004) in which all variables to be imported are estimated in regression equations (chained equations).[5] It can be summarized in the following steps:[6]

- Step 1: Select the P variables $Y_1, Y_2, ..., Y_P$ to be imputed.
- Step 2: Fill the missing values of $Y_1, Y_2, ..., Y_P$ with random selected values which were actually observed.
- Step 3: For each $Y_1, Y_2, ..., Y_P$
  - run a Bayesian regression of the variable to be imputed on a broad set of independent variables, which is chosen from among the HFCS variables without missing values and the variables selected in step 1 (except the one being regressed); the regression sample is restricted to those observations that are not missing in the dependent variable;
  - randomly draw a vector of regression parameters from their posterior distribution;
  - calculate the corresponding predicted values and use them as the imputed values;
  - replace the missing values of the imputed variable with its imputed values.
- Step 4: Repeat step 3 t times. Each time, replace previous imputed values with updated ones obtained from the latest regression. This creates the first imputation sample (or implicate).
- Step 5: Repeat steps 3 and 4 M times independently to obtain M imputation samples.

The basic idea behind this procedure is to impute missing values for each of the P variables with missing values by drawing predictions based on a Bayesian regression model specific to that variable (step 3). To preserve the associations between variables with missing (true) values and variables with *complete* observations, each regression model contains a broad set of independent variables with complete observations.

Furthermore, the procedure is *multivariate* in the sense that the estimation of the missing values is repeated (t times); variables that are being conditioned in each regression are replaced by the observed values or those currently being imputed (step 4). It is important that each regression model also contains a broad set of independent variables with *missing* values in order to preserve the joint distribution of variables with missing values. If t tends to infinity, the imputations of missing values of $Y_1, Y_2, ..., Y_P$ in each cycle are expected to converge to an approximation of a draw from their joint posterior predictive distribution.

In the final step (step 5), the procedure provides multiple imputations of each missing value by repeating steps 3 and 4 M times independently. This is done to take into account the uncertainty of the imputed values when estimating any variances with imputed variables with missing values. The M imputations of the missing values of $Y_1, Y_2, ..., Y_P$ are expected to converge to an approximation of M draws from the joint posterior predictive distribution of the missing values.

Although it is theoretically possible that the sequence of draws based on the regressions above might not converge to a stationary predictive distribution, simulation studies provide evidence that the approach yields estimates that are unbiased (Van Buuren et al., 2006). Furthermore, separate regressions for each variable reflect the data better, given that the HFCS data contain a large number of variables, many of which have bounds, skip patterns, bracketed (i.e. range) responses, interactions or constraints in relation

---

[5] This procedure is also known by several other names, including "stochastic relaxation," "regression switching," "sequential regression," "incompatible MCMC" and "fully conditional specification."

[6] Albacete (2014) provides further technical details on the imputation procedure used for the Austrian Household Survey on Housing Wealth 2008, which is identical to that used for the HFCS.

to other variables. This approach thus makes more sense than specifying a joint distribution for all variables together, as is the case for example in the joint modeling approach.[7]

It should be noted that the HFCS imputation procedure is based on the assumption that the nonresponse probabilities of variables with missing values are only dependent on observed information – never on unobserved information such as the variables with missing values themselves. In the literature this assumption is referred to as ignorability assumption.

Before running through the five steps above, we need to prepare the data and specify all the parameters of our imputation model: e.g. the selection of variables to be imputed, the imputation order, the regression model for each variable, the number of cycles t, the number of imputation samples M, etc. The next section describes how this was done.

5.4 Creating the imputations

5.4.1 Choosing the variables to be imputed

In step 1 of the HFCS imputation procedure, we have to select the variables Y_1, Y_2, ..., Y_P to be imputed. Our strategy is to impute as many variables with missing values as possible, which amounts to around 92% of such variables. This proportion in any case includes all balance sheet variables. The remaining variables with missing values are not imputed with the HFCS imputation procedure due to insufficient variance or a lack of sufficient observations to run a regression.[8]

The imputation of as many variables as possible is intended to minimize the number of cases in which users are forced to conduct a complete-case analysis with HFCS data because the variables they are interested in have not been imputed. Another important reason for adopting this strategy is that we do not want to bias the correlation structure of the data with our imputations. If we were to reject many variables for imputation, we could not use them in the regression models as independent variables with missing values either, and we would thus bias the associations between the unimputed variables with missing values and the imputed ones.

5.4.2 Imputation order

As mentioned in the section on the HFCS imputation procedure, one of the weaknesses of the procedure is that it does not enable us to prove, in theoretical terms, that the sequence of drawn predictions based on the Bayesian regressions converges to a stationary predictive distribution. In practice, however, it has been found that choosing a particular order of Y_1, Y_2, ..., Y_P often aids convergence. Therefore, we order the variables to be imputed by their degree of missingness, starting with the variables with the fewest missing values and ending with those variables that have the most missing values. Variables with the same degree of missingness are imputed in a fixed random order. Head variables are always imputed before their corresponding branch variables. For example, the variable indicating whether a household has a mortgage or not is always imputed before the mortgage amount is imputed, even if the degree of missingness is the same for both variables.

---

[7] See Little and Rubin (2019) for an overview of imputation techniques.

[8] A very small fraction of these variables that could not be imputed with the HFCS imputation procedure were imputed with ad hoc methods such as hotdeck imputation after the HFCS procedure had been completed. This is because their imputation is considered very important as they are used, for example, to calculate important aggregate variables, such as total household income.

5.4.3 Types of regression models

In step 3, we define a regression model for each variable to be imputed. Depending on the type of the variable, we choose from four different types of regression models. For continuous variables, we use a range regression model,[9] because all of our continuous variables are bounded either from above or from below, or both (see section 5.4.6 for more details). For binary variables, we use a logit model; and for ordinal and nominal variables, we use ordered logit and multinomial logit models.[10]

5.4.4 Use of weights in regressions

Generally speaking, there is little debate about the need to use weights for the estimation of descriptive parameters (means, proportions, totals, etc.). There is, however, some debate about the use of weights when fitting regression models to survey data. This issue also arises when fitting the regressions in step 3 of the HFCS imputation procedure. Starting with the second wave, we have used weights as predictors only (section 5.4.7). Put differently, our approach has been not to estimate weighted regressions, in line with the current trend in imputation (see e. g. Frumento et al., 2012). As argued in the literature, multiple imputations are only meant to appropriately predict missing values (and their uncertainty). Units should not be weighted until later, when statements about the population are to be made on the basis of an analysis of the final dataset.

5.4.5 Variable transformations

Before imputing variables with missing values, we transform several of them, as this has proved to be extremely helpful in improving the imputed values of these variables and, hence, in improving the quality of the imputed values in general. Once the imputations are finished, we transform all variables back into their original measure.

One important transformation of continuous variables involves using the natural logarithm. These types of variables usually have a highly skewed distribution; using the logarithm helps to bring the distribution closer to assumption of normality that is necessary for the forecast. Another very helpful transformation for year variables is to impute time periods rather than years. For example, instead of imputing the purchase year of a house, we impute the time elapsed since the house was purchased. In such cases, the logarithmic transformation mentioned above is carried out on the durations and not on the years.

Another transformation used for some variables with values between "0" and "1" is the log-odds transformation $(\log(y/(1-y)))$, for example for the amount of an outstanding consumer loan (HC0801 to HC0803). Instead of imputing these variables individually, the original amount of the consumer loan (HC0601 to HC0603) is imputed as a first step. Additionally, an indicator showing whether the amount outstanding is smaller than the original amount of the loan is imputed, and if so, the outstanding amount is imputed as a percentage of the original amount. This share is imputed as a log-odds transformation,

---

[9] The range regression model is a generalized version of the Tobit model. It is used to account for censoring from below and/or above. See Cameron and Trivedi (2005) for more details.

[10] The nominal variables on the three-digit International Standard Classification of Occupations (ISCO) and the three-digit European statistical classification of economic activities (Nomenclature of Economic Activities – NACE) classifications, which were difficult to estimate with a multinomial logit model because they contain a very large number of categories (74 and 121, respectively), represent the only exceptions. In these two cases, the predictive mean matching (PMM) procedure was used for each missing value to first, predict a value by linear regression and, second, impute the observed value that is closest to the regression-predicted value.

considerably improving the quality of the imputed values. Subsequently, the individual variables (HC0801 to HC0803) are calculated from the original loan amounts and shares.

For categorical variables, two types of transformations may be used. First, some of the nominal variables can be transformed into ordinal variables by reordering categories. This improves the stability of the imputation model, as fewer parameters need to be estimated for ordinal regression models than for multinomial regression models. Second, multiple response variables are transformed into several binary variables by generating one binary variable for each response category ("1" if the category applies, "0" otherwise). This makes it possible to impute more than one response category for the same question per imputation sample.

A transformation that is done for both continuous variables with missing values and categorical variables with missing values involves splitting the original variable into head and branch variables; this is done when there is a certain heterogeneity in the original variable. For example, some loan-length variables have the value "–4," indicating that "The loan has no set length." When imputing such a loan-length variable, it would not make sense to run the regression over these observations together with those variables that do provide a loan-length value. In such cases, the variables are split into two: (1) a binary head variable indicating whether the loan has a set term or not (imputed with a logit regression model), and (2) a continuous branch variable indicating the loan length if the loan has a set term (imputed with range regression).

A further transformation which is carried out both for continuous and categorical variables with missing values is that of individual IDs.[11] Individual variables are modeled and imputed separately for each ID in order to avoid biased imputations (section 5.4.8); this should ensure that household members with the same IDs display relatively homogenous characteristics if they are modeled together. For this reason, respondents are grouped into new individual ID categories created specifically for the imputations prior to imputation. The criteria for this categorization are as follows: All male financially knowledgeable persons (FKPs), all male partners of FKPs that were individual 2 and all other FKPs are classified as individual 1 (ID = 1). All female partners of FKPs that were individual 2 and all women that were individual 1 before their male partners became individual 1 are classified as individual 2 (ID = 2). All other people are ordered and numbered by age in descending order and are numbered starting with ID = 3.

In the case of households with members that engage in farming, we use a special transformation of the variables for the value of the household's business(es) (HD0801 to HD0803) and the variable for the value of the household's main residence (HB0900). Instead of imputing these variables individually, we first impute the sum of these variables and, additionally, the percentage of this sum that is attributable to the farm. Then we calculate the individual variables (HD0801 to HD0803 and HB0900) based on the sum and percentages imputed. The reason for using this transformation is that it considerably improves the imputed values, as some households with members that engage in farming did not state separate values for their main residence and their agricultural business but indicated only the combined value (see section 4.6.2.7 for further details).

5.4.6 Bounds

As mentioned above, we use range regression models to impute continuous variables in step 3 because all such variables are bounded either from above or from below, or both. These bounds are used to avoid the imputation of values that are not defined or that are inconsistent with other variables in the survey. We distinguish between general bounds and individual bounds.

---

[11] In the dataset, financially knowledgeable persons are designated with the ID = 1 by default.

General bounds, which are the same for all households and individuals, are used to avoid imputing values that are not defined or are very unrealistic. Examples of this type of bound include nonnegativity constraints on continuous or count variables (e.g. income or age). For all households the lower bound for these variables is zero. For some continuous variables, we assume that a value above or below a particular general bound cannot occur in practice. As a case in point, the lower bound for the year a loan was taken out (HB1301 to HB1303) is 1945. We assume that no loan that is still outstanding in Austria was taken out, renegotiated or refinanced more than 78 years ago. The use of such empirical bounds helps avoid imputing extreme outliers of these variables without providing biased results. More examples of general bounds include percentage variables (e.g. share of homeownership), where we set the lower bound to zero and the upper bound to 100, or some year variables (e.g. the purchase or inheritance year of the household's main residence), where the upper bound is the year in which the last survey interviews were carried out.

Unlike general bounds, individual bounds take different values depending on each household or individual; they usually ensure consistency with other variables from the same household. Most of the HFCS bounds fall into this category. For example, when imputing the amount spent on food eaten at home, we set the total consumption expenditure estimated by the household as the upper bound. Inversely, when imputing the total estimated consumption expenditure, we set the sum of the amounts spent on food and drink consumed at home and outside of the home as the lower bound. Individual bounds are also used when a household provides a range (either prespecified or individual) in a euro question instead of a specific value. Such ranges are requested if respondents do not provide specific amounts in response to euro questions; they prove very useful for imputation purposes, as they yield valuable and precise information on the missing value from a euro question (see also section 5.2 in connection with table 6).

Individual bounds in the HFCS are, for example, also used when imputing rents (e.g. rent including utilities is used as an upper bound for rent excluding utilities and vice versa), or when imputing several count variables (e.g. the birth year of the oldest household member is used as a lower bound for the year of acquisition of the main residence). If an observation has more than one lower and/or upper bound (e.g. general and individual bounds), we take the lower and/or upper bound that is the most restrictive.

5.4.7 Selecting predictors

As mentioned above, one of the main goals of imputation is to preserve the distribution among variables with missing values and variables with complete observations – and also that among variables with missing values themselves. Therefore, when choosing predictors for the imputation model, it is not sufficient to select the most accurate predictors for each variable to be imputed. Such an approach could bias the correlation structure between the variable to be imputed and the excluded variables. Furthermore, ignoring variables that are determinants of nonresponse for the variable to be imputed makes the ignorability assumption (see section 5.3) less plausible.

Thus, we choose as many predictors as possible (broad conditioning approach). In a large dataset, such as that of the HFCS containing several hundred variables, it is, however, not feasible to include all variables, as this may lead to both multicollinearity problems and computational problems. In line with Van Buuren et al. (1999) and Barceló (2006), we have therefore adopted the following strategy for selecting predictor variables:

1. Include the variables that are determinants of nonresponse. These are necessary to satisfy the ignorability assumption on which our imputation model relies (see section 5.3). Variables included as typical

determinants of nonresponse in the HFCS imputation model are, for example, variables that describe the household (e.g. estimated household income, household size, number of children) and household members (e.g. age, education, sex and employment status of the household's first individual and their partner) as well as stratification variables (e.g. province, municipality size) and information provided by the interviewers (e.g. standard of living, type of neighborhood, building condition, interview atmosphere, etc.). The latter pieces of information (paradata) were extremely important for the imputations, since they provided plausible explanations for item nonresponse for many variables.

2. In addition, include variables that are well suited to predicting and explaining the relevant variable to be imputed. This is the classic criterion for using predictors, and it helps to reduce the statistical uncertainty surrounding the imputations. These predictors are identified by their correlation with the variable to be imputed. For example, when imputing loan variables, we typically use the original loan amount (as mentioned above), the repaid loan amount or principal outstanding as predictors because, in most regressions, these variables can explain a considerable amount of variance. When imputing the market value of various types of real estate property, we usually include the purchase value and the length of time (in years) for which the household has already owned the respective property. When imputing loan variables, we typically (as described above) use the original amount, the loan repayment amount or the loan amount outstanding. Usually, these variables are connected logically (e.g. the outstanding principal is the original loan amount minus the sum of all loan repayments). However, in the course of imputation, it is not possible to preserve all of these logical connections, in particular if several of these variables are being imputed simultaneously.

3. Remove the aforementioned predictor variables that have too many missing values in the subsample of missing observations of the variable to be imputed and substitute them with more complete predictors of these predictors. As a rule of thumb, predictors where the percentage of observed cases within this subsample is below 50% are removed and replaced by more complete predictors. This criterion helps to make the imputations more robust. Typical predictors of predictors include essential household characteristics, such as household size, the number of children, region, age, as well as the employment and marital status of the first individual.

4. Include all variables that appear in the models that will be applied to the data after imputation. In other words, consider which economic theories might be tested based on the data and include those variables as predictors that are expected, according to these theories, to influence or explain the variable to be imputed. Failure to do so will tend to bias the results of potential data users when testing the hypothesis of one particular model. For example, the HFCS data provide detailed information on different components of households' wealth, e.g. real assets or financial assets. This information is used for the analysis of wealth effects on consumption. Therefore, we use these variables both for the imputation of consumption expenditure and for the imputation of wealth variables.

Obviously, many variables in the survey – for example, the income, age or education of the first individual – fulfill more than one criterion for selecting predictors.

We also include the final survey weights in all regression models (see the discussion in section 5.4.4) and an interaction term, as well as a main effect dummy for each of the abovementioned predictor variables that were only asked from a subsample of the households asked about the variable to be imputed. For example, suppose that we want to impute a household's consumption expenditure using the mortgage amount as one of our predictors. While every household in the sample was asked about consumption expenditure, not all of them were asked about mortgage amounts. If, for those households that do not

have a mortgage, we just set the mortgage amount to zero (corresponds to an interaction term), the estimates would be biased, because the information on whether a household has a mortgage or not would be omitted. This information should thus be additionally included as a main effect dummy in the regression model. But again, not all households were asked whether they have a mortgage, just homeowners. Thus, we should also include a homeowner dummy in the regression.

Finally, the number of predictors is restricted by the size of the subsample for which the regression is estimated. In cases where the subsample size is smaller than the number of predictors selected according to the above strategy, we use the Akaike information criterion to choose the subset of predictors which best fits the data, ensuring that, if possible, each of the above four predictor categories is represented in each regression equation. Typically, the number of predictors used for each regression model is around 20% of the number of observations for the variable to be imputed. More details on the specification of subsamples can be found in the next section.

5.4.8 Specification of subsamples

Each regression in step 3 is estimated over a subsample consisting of all households and individuals that were asked the question pertinent to the variable to be imputed. For example, if a household has two mortgages and we want to impute the outstanding amount of the second mortgage, then we impute this missing value by regressing over the subsample of households that have at least two mortgages. If we also included the households that only have one mortgage when imputing the second mortgage amounts, we would ignore systematic differences between the first and second mortgages. For example, we would ignore the fact that the outstanding amount of the first mortgage is always higher than the second one, because mortgages are ordered by importance, which would introduce a bias to our estimates.[12]

A further example is the imputation of individual variables. These are also only regressed over the subsample of people that share the same ID. To ensure the homogeneity of people with the same IDs, respondents are grouped into new ID categories created specifically for the imputation (see section 5.4.5) that then form the mentioned subsamples. When imputing question by question, as we do, the bias will be very small, though at the cost of precision because, consequently, the subsample sizes are often small.

5.4.9 Number of cycles

In step 4, the number of cycles (or iterations) t determines how often step 3 is repeated. As t tends to infinity, the imputed values should converge to a draw from the joint posterior predictive distribution of the variables with missing values. However, according to Van Buuren et al. (1999), in practice, convergence in these models usually occurs very quickly during the first few iterations. Given the large computational effort required for the HFCS imputation model, we set the iteration number for the HFCS imputation model at $t = 10$, which is in line with other similar surveys, like the SCF (Kennickell, 1998) and the EFF (Barceló, 2006).

Typically, we check convergence graphically by plotting the mean of the imputed values against the iteration number t. Convergence is judged to have occurred as soon as the pattern of the imputed means becomes random and a definite trend can no longer be observed.

---

[12] Even if, in such cases, we could introduce a large number of interaction terms to our model to reduce the bias, there might still be unobserved differences between the two groups.

Starting with the second wave of the HFCS, we additionally examine the convergence of selected variables using the Gelman-Rubin convergence diagnostic, which is used very frequently in the literature (for more details, see e.g. Cowles and Carlin, 1996). According to this diagnostic, convergence of a variable is reached when the variance of an estimate of this variable (e.g. the mean, median or other percentiles) is relatively small within chains of multiple imputation samples compared to the variance of the same estimate between cycles.[13] All variables examined in the fifth wave of the HFCS meet this criterion for almost all their estimates.[14]

Of course, such tests (just like any other diagnostic test to assess chain convergence) can never confirm the existence of convergence (see section 5.3). But they are useful for pointing out weaknesses of the imputation model or other unusual results that could indicate nonconvergence.

5.4.10 Number of imputation samples

In the last step (step 5), we choose the number of realizations $m = 1,2,..., M$ that we want to have from the joint posterior predictive distribution of the missing data or, put more simply, the number of samples to be generated through multiple imputation. Setting M too low leads to standard errors of estimates that are too low and to p-values that are too low. However, Schafer and Olsen (1998) show that the gains in efficiency of an estimate rapidly diminish after the first few M imputation samples. They argue that good inferences can already be made with $M = 3$ to $M = 5$. In line with the international standards set by the ECB and other similar surveys (like the SCF or EFF), we set the number of imputations at $M = 5$.

5.5 Selected results

After imputation, the HFCS dataset is five times bigger than before, because it consists of $M = 5$ multiple imputation samples (also referred to as "implicates"). Table 8 provides first insights into the imputation output. It shows the weighted means of selected euro variables in both the multiple imputation samples and the original unimputed sample.

For several variables the means are, on average, higher after imputation than before imputation. If imputations are close to the true values, the result suggests that households that do not respond to the relevant variables tend to be households with higher (unobserved) amounts in these variables. For example, the mean value of the first gift/inheritance (without main residence) is EUR 98,565 before imputation. After the respective imputations, it increases to EUR 131,170 in $m = 1$, EUR 125,522 in $m = 2$, EUR 132,980 in $m = 3$, EUR 121,765 in $m = 4$, and EUR 138,726 in $m = 5$. Note that the ECB changed the recording of inherited HMR into the loop. Thus, on average the imputations increase the mean value of the first gift/inheritance from EUR 98,565 to EUR 130,033, i.e. by 32%. Additionally, about 40% of the values imputed in this context are based on range responses by households, which suggests that households with more valuable inheritances tend to respond to this question less often than households with less valuable inheritances. Further large increases in comparison to the unimputed sample occur when imputing properties other than the main residence. Households' range responses again play an important part here, as they provide valuable and often very precise information for the imputations (see also table 6).

---

[13] The Gelman-Rubin diagnostic is the *square root of [(t–1)/t) + (BV/WV)]*, with BV denoting the between-chain variance and WV the within-chain variance. If the Gelman-Rubin values are below 1.2, they are usually considered to denote convergence.

[14] Exception: mean and 5th percentile of *HD1510*. The following variables were tested: *HB0900, HD1110, HD1210, HD1510, HB1701, HB2801, HB4400, HI0100, HI0200* and *HI0310*.

However, for other variables, the mean does not change significantly, or even decreases in some cases. For example, the mean amount spent on food eaten at home does not change significantly after imputation, due to the low item nonresponse rate of this variable (see table 6). The mean gross income from financial investments is even lower after imputation than before imputation, which suggests households receiving a substantial amount of income in this category are more likely to know it.

Finally, table 8 also shows that the uncertainty of imputations can vary a lot depending on the variables. For some variables (e.g. other property mortgage 1), the means show a relatively high variance among the five multiple imputation samples, signaling the uncertainty of the imputed values due to the lower number of observations for these variables. For other variables (e.g. gross income from unemployment benefits or the monthly amount paid as rent) the mean values show a relatively low variance among the five multiple imputation samples, which in turn signals a higher precision of the imputed values. Had we conducted a single imputation of the variables – with only one imputation sample – instead of multiple imputations, the variance of the estimates would be too low, since the uncertainty behind the imputed values would be disregarded, and they would thus be treated like true values. The variance among the five multiple imputation samples is within the range of the previous waves.

Table 8

Means for selected variables before and after multiple imputation (weighted)

	Mean before imputation	Multiple imputation sample means				
	m=0	m=1	m=2	m=3	m=4	m=5
	EUR					
Value of main residence[1]	444,962	443,604	444,302	443,584	441,188	442,558
HMR mortgage 1: amount still owed	92,737	90,265	96,340	94,119	95,958	100,228
Monthly amount paid as rent	509	474	478	471	472	472
Other property 1: current value	326,999	335,386	352,661	351,911	344,444	340,742
Other property mortgage 1: amount still owed	80,899	81,019	89,696	75,370	80,544	74,607
Value of sight accounts	5,053	5,071	5,066	5,079	5,217	5,197
Value of saving accounts	28,816	29,998	30,584	30,310	30,594	30,121
Value of publicly traded shares	55,888	47,509	45,715	45,641	51,465	44,522
Gross employee income (person 1)	36,674	36,344	36,211	36,342	36,480	36,179
Gross income from unemployment benefits (person 1)	8,432	8,222	8,218	8,155	8,117	8,258
Gross income from financial investments	406	326	327	339	338	317
Gift/inheritance 1: value	98,565	131,170	125,522	132,980	121,765	138,726
Amount spent on food at home	452	453	454	453	454	453

Source: HFCS Austria 2023, OeNB.

[1] Based on the HB0900 variable.

Note: All means are estimated over the observations "Household has item = yes." The number of these observations may vary across the different imputation samples m if we impute whether households have the relevant item or not. HMR = household main residence.

5.6 Concluding remarks

We have shown that imputation is necessary for analyzing the HFCS dataset because, compared with complete-case analysis, it decreases the nonresponse bias of estimates when complete observations differ systematically from incomplete ones. It also decreases the loss of information in analyses because no observations need to be deleted. We chose a multiple imputation with chained equations to create five multiple imputation samples. For information on analyzing multiply imputed data in Stata and R, please see the HFCS user guide (chapter 9).

6 Sampling

6.1 Introduction

The sampling design for the fifth wave of the HFCS in Austria was specifically developed by the OeNB in collaboration with the survey company IFES (Institut für empirische Sozialforschung GmbH), which also executed the survey. Sampling is understood as the selection of a set of units (i.e. a sample) from the whole population on the basis of which conclusions can be derived about the behavior of the whole population. Thus, the units of the sample should be representative of the whole population; in other words, an analysis of the sample (using appropriate weights) is expected to lead to the same estimates as an analysis of the whole population. Another criterion of major importance for the HFCS – coverage of households in all geographic regions – is achieved through stratified sampling, i.e. by dividing a country into smaller units from which the sampling units are drawn. Although some degree of statistical uncertainty cannot be ruled out, sampling – together with imputation and weighting – serves to produce best unbiased estimates (and confidence intervals). It further keeps uncertainty as low as possible taking restrictions like cost, time and practicability into account. Therefore, every survey is highly dependent on the quality of its sampling design.

This chapter describes the sampling procedure for the HFCS in Austria and is structured as follows: First, we define the target population (section 6.2) and provide a short overview of the sampling design in box 1. This part is followed by a description of the required external data on geography and population (section 6.3). Next, we detail the stratification process (section 6.4) and the two stages of drawing the survey's sample population (section 6.5), which form the main part of the sampling procedure. Section 6.6 completes the chapter with some concluding remarks.

6.2 Target population and sampling frame

The first step in determining the sampling procedure is to define the target population of the survey. The survey is intended to cover all households living permanently in Austria, independent of citizenship and/or residence status. In line with the common ECB definition, a household in the HFCS is defined as "a person living alone or a group of people who live together in the same private dwelling and share expenditures, including the joint provision of the essentials of living. Employees of other residents (i.e. live-in domestic servants, au-pairs, etc.) and roommates without other family or partnership attachments to household members (e.g. resident boarders, lodgers, tenants, visitors, etc.) are considered separate households."[1]

More specifically, according to the ECB's definition the following persons are to be regarded as household members if they share household expenses (which includes both benefiting from and contributing to their coverage):

1. persons usually resident, related to other members
2. persons usually resident, not related to other members
3. persons usually resident, but temporarily absent from dwelling (for reasons of holiday travel, work, education or similar)
4. children of household being educated away from home
5. persons absent for long periods, but having household ties: persons working away from home
6. persons temporarily absent but having household ties: persons in hospital, nursing home, boarding school or other institution.

---

[1] See ECB (2013a), p. 80f. and ECB (2016), p. 11f.

Sampling in the HFCS in Austria

The HFCS in Austria is based on a stratified, two-stage cluster sampling design:

"Stratified sampling" ensures that the data collection units – i.e. households for our purposes – are drawn from all parts of the population. Stratification in the Austrian HFCS was carried out geographically (based on NUTS-3 regions[2]) and by municipality size.

Figure B1

Overview of two-stage sampling

Source: Statistics Austria, HFCS Austria 2023, OeNB.

"Two-stage cluster sampling" refers to a process where primary sampling units (PSUs) are first selected from each geographical unit (i.e. stratum) and secondary sampling units (SSUs) are then drawn from within each selected PSU. The two-stage sampling design of the Austrian HFCS (see the infographic below) entails, first, selecting a random sample of enumeration districts (the smallest geographical unit for which statistical data are available) from each stratum and, second, selecting a random sample of households (postal addresses) from each sampled enumeration district. As in the previous waves of the HFCS in Austria,[3] the probability of being drawn during this first stage is proportional to the number of households in the respective PSUs.[4] The households constitute the secondary sampling units (SSUs) and are selected at random from within a drawn PSU. The two-stage cluster design reduces costs due to relatively small distances between the 12 households (8 households in strata with over 50,000 inhabitants) selected within each PSU, while it ensures a sufficient number of PSUs within the individual stratum.

This guarantees that households from every single stratum are invited to take part in the survey. In total, the gross sample of the fifth wave of the Austrian HFCS consists of 192 strata, 689 unique PSUs and 7,636 households.

---

[2] See https://www.statistik.at/en/services/tools/services/regional-information/regional-divisions; accessed on July 29, 2025. Austria is divided into 35 NUTS-3 regions. These regions typically consist of several neighboring political districts or correspond to urban areas including the capital cities of the provinces.

[3] In the first wave this was not the case.

[4] In accordance with the literature (see Williams, 2014; or Valliant et al., 2013), if the calculation of this probability was "greater than one," then these PSUs were put into the sample with probability one.

In the case of the HFCS, the target population does not include households that are institutionalized, such as households living in

- homes for elderly people,
- military compounds,
- monasteries,
- prisons, and
- boarding schools.

Additionally, the Austrian HFCS does not cover homeless people. People without a residence were not reached with the survey, as sampling was based on fixed dwellings (see section 6.3). At the same time, the HFCS in Austria is not limited to households officially registered at their main residences.

In order to draw a sample from this target population, we would need a complete list of households in Austria. As such a list does not exist, we use a complete list of postal addresses in Austria as our sampling frame. These external data, explained in more detail below, provide the best possible sampling frame in the sense that (almost) all households in Austria appear in the data (and appear only once) and that the data are highly up to date.

6.3 Background – the (external) datasets used

Given the definition of the target population, geographical data, as well as data on households in Austria are needed. A representative draw from the sample requires the target population to be correctly represented by the sampling frame. The frame data are perfect "if every element appears on the list separately, once, only once and nothing else appears on the list" (Kish, 1995, p. 53). In practice, it is not possible to achieve this theoretical optimum. The Austrian HFCS has been designed to meet this goal subject to the constraints of the data sources available. The following sections give more details on the data which provided the basis for the sampling used in the HFCS in Austria.

For the HFCS in Austria, we relied on two different sources: We used data from Statistics Austria for the purpose of stratification and for selecting a random sample of PSUs (primary sampling units; in Austria those are the enumeration districts) and we used post office data to draw the households, the actual SSUs (secondary sampling units), at random. The advantage of the post office data is that they are up to date and that the data fit the HFCS definition of households.

6.3.1 Statistics Austria

We used information about the geographical structure of Austria, i.e. data on the NUTS-3 regions, and the enumeration districts (PSUs) from Statistics Austria.[5] These enumeration districts are the smallest territorial units in Austria for which basic data characteristics are collected by Statistics Austria by default (each enumeration district contains around 463 dwellings on average).[6]

In addition, we relied on the municipality directory of 2022 (to categorize by municipality size). Unlike in the early waves, we did not use population data from Statistics Austria's register-based census from 2011 to determine the number of sampling units to be drawn within each stratum. Instead, we relied on post office data, which also contained the population (households) of each stratum (see section 6.3.2 below).

---

[5] See also https://www.statistik.at/en/services/tools/services/regional-information/regional-divisions (accessed on August 8, 2025).

[6] The estimated number of households in Austria according to the HFCS definition (4.1 million) divided by the number of enumeration districts (8,842) yields 463.2.

6.3.2 Austrian Post

Once the appropriate primary sampling units have been randomly selected, information on the households is needed to complete the sample selection. The dataset of choice for the purpose of the HFCS was a dataset of postal addresses for sale from the Austrian Post, based on the assumption that the number of households living in each building corresponds to the number of postal addresses. Specifically, we used a commercial product called "Adress.Certified" developed by the Austrian Post. This address register contains information about individual buildings (including street name, house number and whether the building is used privately or commercially). It can be purchased in combination with a product called "DATA.DOOR," which is a directory of post office-certified address codes (shortened to PAC), i.e. a directory of all addresses in Austria that mail can be delivered to. This information is available in disaggregated form. In Austria, there are about 4,621,000 private mail delivery points. Addresses identified by the Austrian Post as vacation homes have already been excluded.

Thus, our starting point was some 4.62 million private mail addresses. Very few remaining commercial addresses and ineligible addresses had to be removed after the first contact by the interviewer (e.g. if the interviewer arriving at a given address reported the address as incorrect or found it to house a commercial building) and were given weights of zero, since they do not belong to the target population (see chapters 4 and 7). Moreover, households at secondary residences whose main residence address was identifiable as such were also excluded from the dataset of the sampling frame or given a weight of zero to ensure that every household was included only once in the list of post office-certified address codes. After removing these commercial addresses and ineligible addresses, the total of all weights comes to roughly 4.095 million, which means that Austria has an estimated 4.1 million households.

The post office data we used do not reflect whether a given address is a household's main residence or not or whether the residence at that address is registered in the Austrian residence registry (Zentrales Melderegister). Yet they provide a realistic picture of households and thus meet the HFCS requirement of reflecting actual living situations. Unlike other data sources (e.g. EU SILC), the post office data cover households at addresses that are registered as a secondary home or that are not registered at all, but fulfill the HFCS definition of a household. They have thus been included in the sampling frame because they have a post-certified address code.[7]

6.3.3 Profile.Address and IFES

To identify the names of the households that correspond to the selected postal addresses – information that is not evident from the datasets described so far – the survey company, IFES, relied on its databases or obtained the corresponding addresses from a commercial provider called "Profile.Address."

This information was needed in the contact phase when households received individualized letters of invitation to participate in the survey.[8]

---

[7] The post-certified addresses for some 4.6 million households compare with about 4.1 million household addresses documented by other sources (such as the microcensus based on the residence registry; see also section 7.2.4).

[8] For more details on the contact strategy, see section 3.4.

6.4 Stratification and sample size

6.4.1 Stratification

The Austrian HFCS essentially used two indicators for stratification, the first one being the 35 NUTS-3 regions (see figure 4).

With the exception of the capital city Vienna, each NUTS-3 region was divided further into the following eight municipality population size categories. The data from Data.Door from the Austrian Post contain information that serves to map each address to a province, municipality and enumeration district as they are defined by Statistics Austria.

- Up to 2,000 inhabitants
- 2,001 – 3,000 inhabitants
- 3,001 – 5,000 inhabitants
- 5,001 – 10,000 inhabitants
- 10,001 – 20,000 inhabitants
- 20,001 – 50,000 inhabitants
- 50,001 – 1 million inhabitants
- Over 1 million inhabitants

The large "50,001 – 1 million inhabitants" category essentially contains just the provincial capitals. Vienna is a special case as it is the capital city and the only city in Austria with more than 1 million inhabitants; it was subdivided into its 23 districts.

This very fine stratification yielded 199 strata. Where the number or proportional share of households per stratum was too small to allow the selection of enumeration districts, individual strata were merged

Figure 4

NUTS-3 regions

Source: Statistics Austria.

with neighboring strata to increase the share of households and thus ensure the selection of at least one PSU from each stratum. This exercise left the HFCS with 192 strata for sampling covering all households in Austria. The distribution of strata across provinces and municipality size categories can be seen in table 9.

Each stratum contained about 44 PSUs on average, which in turn contained around 486 households[9] on average.

Allocation of strata within the sample

	Municipality size (number of inhabitants)[1]								
	Up to 2,000	2,001– 3,000	3,001– 5,000	5,001– 10,000	10,001– 20,000	20,001– 50,000	50,001– 1 million	Over 1 million	Total
Vienna	0	0	0	0	0	0	0	23	23
Lower Austria	7	7	7	7	6	5	1	0	40
Burgenland	3	2	3	2	1	0	0	0	11
Styria	6	5	6	6	5	1	1	0	30
Carinthia	3	3	3	3	2	1	1	0	16
Upper Austria	5	5	5	5	3	2	1	0	26
Salzburg	3	2	3	2	2	1	1	0	14
Tyrol	5	5	3	3	3	0	1	0	20
Vorarlberg	2	2	2	2	2	1	1	0	12
Total	**34**	**31**	**32**	**30**	**24**	**11**	**7**	**23**	**192**

Source: Statistics Austria (municipality directory 2022).

[1] Municipality size accounts for municipality mergers up to and including 2022.

6.4.2 Sample size

The variance of estimates based on the underlying data will be smaller the larger the sample is. At the same time, the cost of data collection increases with sample size. Therefore, a balance has to be found in order to yield reasonably precise estimates whilst taking into account the given budget constraints. Furthermore, given the focus of the survey and the analyses that the HFCS is to carry out, the HFCS should produce enough observations to allow for an analysis of subpopulations (e.g. indebted households, which are only a (small) fraction of the sampling frame) and provide some insight into the regional differences within Austria. We know from previous OeNB surveys (e.g. the four previous waves of the HFCS in Austria in 2010, 2014, 2017 and 2021 as well as the Household Survey on Housing Wealth in 2008) that at least 2,000 households need to be successfully interviewed and that the unit nonresponse rate can be expected to reach some 40% to 50% (with expected differences between Vienna and the rest of Austria).[10] With some leeway for extraordinary circumstances, the HFCS survey was designed to yield a sample of 3,000 successfully interviewed households and a participation rate of about 37% in big municipalities (more than 50,000 inhabitants) and an average of around 41% in the rest of Austria. These participation rates are estimates based on the experience of past surveys. The participation rates recorded in the previous waves in a stratum were used to determine the exact number of households to be drawn in the respective stratum in the fifth wave.

[9] The difference between this figure and the 463 households per PSU cited earlier is the result of aggregation (see also section 6.5.1). Similar to the data above we used the roughly 4.1 million households according to the HFCS definition in the 8,429 enumeration districts resulting from the aggregation to calculate this average.

[10] The experience of regional differences between participation rates in past waves of the HFCS in Austria were taken into account in the sample design for the fifth wave.

Table 10

Identification of the number of primary sampling units (PSUs) to be drawn

	% of households	Target sample	Gross sample	Number of households per PSU (enumeration district)	Number of PSUs to be drawn
Vienna	25.1	752	2032	8	254
Lower Austria	18.7	562	1376	8 / 12	117
Burgenland	3.3	98	240	12	20
Styria	13.8	415	1048	8 / 12	101
Carinthia	6.1	182	480	8 / 12	47
Upper Austria	14.9	448	1092	8 / 12	102
Salzburg	6.0	179	440	8 / 12	43
Tyrol	8.0	240	604	8 / 12	56
Vorarlberg	4.1	124	324	12	29
Total	**100**	**3.000**	**7.636**		**769**

Source: Austrian Post, data.door October 2022, HFCS Austria 2023, OeNB.

The targeted net sample of n = 3,000 was divided between the nine provinces, based on their share of private addresses as collected by the Austrian Post (table 10, column 1). These figures, which corresponded to the targeted number of secondary sampling units (SSUs, column 2), were subsequently translated into gross samples of SSUs based on the estimated participation rates (column 3). Due to the shorter distances between buildings, 8 households were selected in Vienna and in strata with more than 50,000 inhabitants, whilst this number was 12 in the rest of Austria (column 4). The number of PSUs to be drawn in each province was calculated on this basis (column 5).

In total, the Austrian HFCS sample design produced 769 (689 unique) PSUs across all strata and a gross sample size of 7,636 households that were invited to participate in the HFCS (see box 2). Drawing the PSUs was done with replacement sampling, which entails some PSUs being drawn multiple times (see section 6.5.1). Drawing possible substitute addresses was excluded from the HFCS to begin with to ensure that all households from the gross sample would be interviewed with the same commitment so as to prevent data distortions (see also section 4.4.1).

6.5 The two stages of the random draw

The Austrian HFCS is based on a stratified, two-stage cluster sampling design:
- stage one: random draw of PSUs (enumeration districts) from each stratum
- stage two: random draw of a predefined number of households (postal addresses) from each PSU

6.5.1 First stage

We chose the smallest territorial units in Austria, the so-called enumeration districts (of which there are 8,842), as the PSUs for the Austrian HFCS. On average, an enumeration district contains 463 households, but there are also PSUs with only a few households. Such units were aggregated with neighboring units to ensure that each PSU contains at least 50 households and that at least one PSU can be chosen per stratum. This aggregation process narrowed the number of PSUs down to 8,429, which then had 486 households on average. The description above shows that the number of PSUs to be drawn in each province is determined a priori by the chosen sample size and stratification. To translate the numbers allocated to each region (table 10) into the desired number of PSUs within a given stratum, the total number of

PSUs in the respective province was distributed proportionally according to the number of households in the respective stratum. For example, the 102 PSUs to be drawn in Upper Austria (table 10) were divided up among the 26 strata in this province according to its population share.

Table 11

Matching of Statistics Austria data with post office and commercial data (fictitious example)

First stage		Second stage				
Statistics Austria		Austrian Post				Profile.Address/IFES
Municipality code (1)	Enumeration district (2)	Postal code (3)	Street (4)	House number (5)	Mail delivery point (PAC) (6)	Name of household (7)
90101	90101001	XXXX	Sample street	6	101255765	John Doe
90101	90101001	XXXX	Sample street	6	101255766	Jane Doe
90101	90101002	XXXX	Sample street	9	101255767	John Doe
90101	90101001	XXXX	Sample street	10	101255768	Jane Doe

Source: Statistics Austria, Austrian Post, Profile.Address/IFES.

After determining how many PSUs were to be drawn per stratum, the PSUs – unlike during the first wave of the HFCS in Austria[11], but as in the second, third and fourth wave – were drawn proportionally to their size (measured in terms of the number of households in a PSU).[12] The purpose of this approach is to reduce the standard errors of estimates by reducing sample design weight variance (see also section 7.2.2). Likewise, it ensures that every household within a stratum has the same probability of being drawn in the gross sample of the HFCS. PSUs were drawn with replacement, meaning that a PSU can be drawn multiple times. This meant that a total of 769 PSUs were drawn in the fifth wave of the HFCS in Austria, only 689 of which were unique.

6.5.2 Second stage

With 769 (689 unique) PSUs having been randomly drawn, we turn to the second stage in which households are selected.

Mail delivery points were randomly selected from each PSU drawn, with 8 being chosen in Vienna or in strata with more than 50,000 inhabitants and 12 being chosen in all other strata. In this process, every household in a given PSU has an equal probability of being selected in the sample, which is measured as a ratio of 1 to the number of households in that PSU. This procedure resulted in a gross sample of 7,636 households in Austria.

6.5.3 Practical implementation

Table 11 illustrates how the data from the second stage of the sampling was used, after the PSUs had been chosen in the first stage (column 2): Austrian Post data (column 6) were used to determine the appropriate mail delivery point, which gives the address, but not the holder of this address. To identify the

[11] During the first wave, the probability of drawing a PSU within a stratum was identical for every PSU.

[12] Mathematically, the probability of an enumeration district being drawn in a given stratum can be expressed as the number of households in a given PSU divided by the total number of households in this stratum times the number of enumeration districts drawn.

name of the household corresponding to the selected postal address, the survey company, IFES, used its own databases or, where necessary, bought the corresponding name from the company "Profile.Address" (column 7).

Since the first contact with a household is very important for a successful interview, every household selected for the HFCS survey received an individualized letter signed by the Governor of the OeNB. This letter contained information on the survey and an invitation to take part (see section 3.5.1).[13]

6.6 Concluding remarks

This chapter provides information on the sampling design as developed for the fifth wave of the HFCS in Austria on the basis of the designs for the first four waves. Improvements reflect the experience of past waves. As described, the survey is based on stratified, two-stage cluster random sampling, consisting of a random draw (proportional to the number of households) of primary sampling units (PSUs; here, enumeration districts) from each stratum plus a random draw (with an identical probability of being drawn within a given PSU) of households (postal addresses as available from an Austrian Post database) from the selected PSUs.

- The sampling method used for the HFCS has a number of advantages, with the following aspects being particularly important:
- The fact that the probability of drawing a PSU was proportional to its size in terms of the number of households ensured a better efficiency of sampling design as compared to a design in which all PSUs could be drawn with the same probability by reducing the variance of the design weights.
- As sampling does not differentiate between main residences and second homes (as recorded in the residence registry), all households that correspond to the HFCS household definition have a positive probability of being selected.
- The very fine stratification structure ensures that all segments of the Austrian population are represented in the survey.

At the same time, given the topics covered by the HFCS it would be desirable to oversample certain groups of the population, such as wealthy households, to improve the efficiency of estimates for these subgroups. However, management decided against the implementation of such an oversampling scheme.

---

[13] See the online appendix for the invitation letter.

7 Construction of survey weights

7.1 Introduction

Survey weights are usually constructed for two reasons: first, to make the sample representative of the target population and second, to reduce sampling variance.
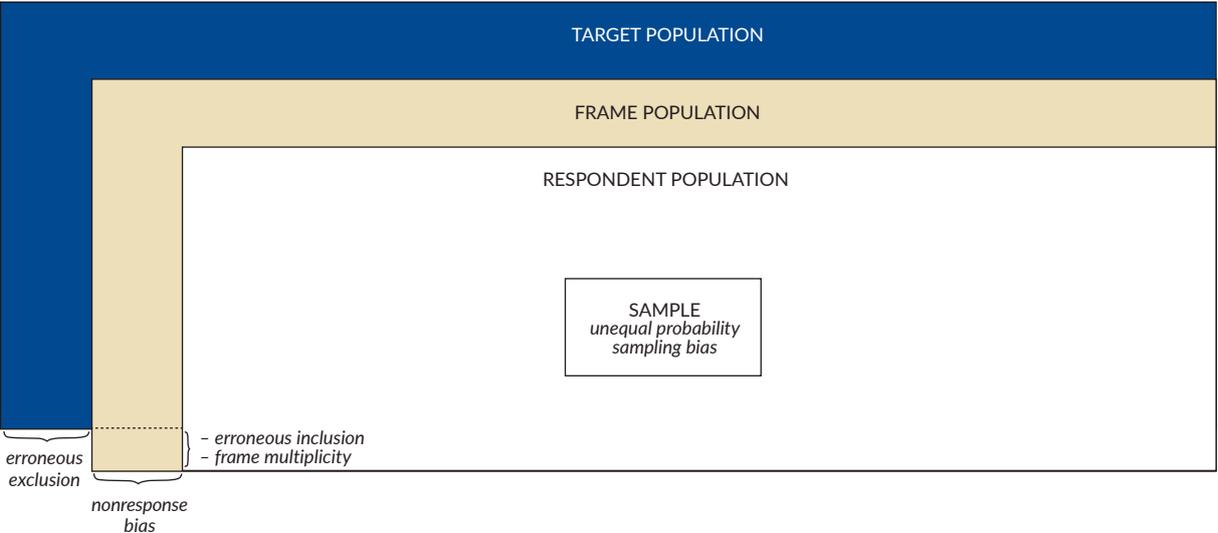
The target population of the HFCS consists of all households in Austria, with a household being defined as an individual or a group of people who live together in the same private dwelling and share expenses.[1] However, the sample may contain several types of biases that may cause a misrepresentation of this target population: unequal probability sampling bias, frame bias and nonresponse bias (see figure 5).

As mentioned above, the unequal probability sampling bias is due to the fact that not every household has the same probability of being selected into the sample, reflecting the fact of oversampling of households in urban areas (like Vienna) in the HFCS sample, which is used to address the known problem of the relatively low survey participation propensity of urban households. To correct these misrepresentations, we constructed design weights, which will be explained in section 7.2.2. Further details about the HFCS sampling design can be found in chapter 6.

Imperfections in the survey frame from which the sample is drawn can lead to frame bias. In the HFCS, the sampling frame is a list of all personal postal addresses in Austria (see chapter 6). Erroneous exclusion of households could imply an imperfection with respect to the target population. In other words, there is the possibility that households without a postal address, for example, one-person households living together in residential communities and sharing an address that contains only one of these households, were excluded. These types of households would then be underrepresented. Another imperfection of the frame could be caused by erroneous inclusion, that is, the inclusion of addresses not belonging to households,

Figure 5

Misrepresentation of the target population in the sample

Source: Adapted from Biemer and Christ (2008).

---

[1] Some special types of households, like those living in care residences (retirees, people in need of care), prisoners, etc., are excluded from this definition. For more details on the definition of the target population, see chapter 6.

for example, those of companies or households in care residences.[2] Finally, there is a third type of imperfection called frame multiplicity, which means that households may be duplicated because they have two (or more) addresses, for example multiple domiciles of commuters. Depending on its type, the frame bias can be reduced by using design weights[3] (to address erroneous inclusion and frame multiplicity) or poststratification weights (to address erroneous exclusion). We explain the construction of these weights in more detail in sections 7.2.2 and 7.2.4.

The nonresponse bias is caused by the fact that only a subset of the households included in the gross sample is willing to participate in the survey. Certain groups of households have a lower probability of participating in the HFCS than other groups – a phenomenon widely corroborated in literature (see e.g. Kennickell and McManus, 1993). Thus, estimates for the sampling frame would be biased with respect to these group characteristics, even though they are unbiased for the participating population. Using nonresponse weights can correct this bias (section 7.2.3).

Furthermore, as mentioned above, survey weights can help to reduce sampling variance, and, hence, to increase the precision of the estimators. Ideally, the precision of the estimators should be improved by stratification prior to sampling. However, some variables (e.g. household size) that would have been very good for stratification and, thus, for improving the precision of the estimators, were not available until after the sample had been drawn and the sample households had been contacted. Some of the gain in precision that would have been possible by using these variables for stratification can be achieved by using these variables for poststratification. These poststratification weights were also utilized for correcting erroneous exclusion (see section 7.2.4).[4]

The construction of survey weights is very important for the HFCS. The following sections will explain how design, nonresponse and poststratification weights were constructed and how the final set of survey weights was derived from these weights. Finally, we will present some descriptive results that take these weights into account.

7.2 Construction of survey weights

7.2.1 Weight components

We aim to construct a final survey weight w_i for every household i that is relatively small for households that are overrepresented in the sample compared to the target population and relatively large for households that are underrepresented. However, as already mentioned in the introduction, households may be misrepresented with respect to the target population for various reasons. Therefore, a specific adjustment using weights is required for every type of misrepresentation. In the HFCS, three types of weights are used: design weights w_{Di}, nonresponse weights w_{NRi}, and poststratification weights w_{PSi}. The product of these three weights yields the final survey weight w_i:

$$w_i = w_{Di} \cdot w_{NRi} \cdot w_{PSi}$$

---

[2] Although addresses of companies and of households in care residences were removed from the sampling frame, some such addresses may still be erroneously included.

[3] Sometimes referred to as noncoverage bias.

[4] Poststratification weights can, moreover, correct a third type of sample-specific bias: the target population may be accidentally misrepresented by the specific households drawn into the sample.

Although some HFCS variables are asked at the individual level rather than the household level, no weights were constructed for individuals because the main focus of the survey is the household level.

7.2.2 Design weights

Design weights help reduce the unequal probability sampling bias, as well as rectify erroneous inclusion and frame multiplicities. In the HFCS, we compute the design weights on the basis of two-stage cluster sampling and the selection probabilities of the primary sampling units (PSUs) and the secondary sampling units (SSUs). In the first stage, the smallest territorial units, the so-called enumeration districts (PSUs), are drawn; then in the second stage, the households (SSUs) within these enumeration districts are drawn (see section 6). The probability that the i^{th} household in the j^{th} enumeration district is selected into the sample is the product of the selection probability for the enumeration district and the selection probability for the household, under the condition that the household's enumeration district is selected. The inverse of this product is the preliminary design weight. The calculation of the design weight mirrors the two steps of the sampling procedure:

Step 1: Calculate the probability that a certain PSU is selected. As described in chapter 6, this sampling probability is defined depending on the relative number of households in a PSU. The probability that PSU j will be selected in stratum h is

$$PSU(h, j) = \frac{M_{hj} m_h}{N_h},$$

where M_{hj} represents the number of households in this enumeration district (h,j), m_h the number of PSUs to be drawn in this stratum, and N_h the number of households in this stratum.

Step 2: Calculate the probability that a SSU is selected. Under the condition that a PSU is chosen, each household in this enumeration district has the same probability of being chosen. Thus, the probability of being selected is given by

$$\frac{m_{hj}}{M_{hj}},$$

where m_{hj} is the number of households to be drawn in the PSU (i.e. 8 in a stratum with a population of over 50,000 and 12 in the rest of Austria). As above, M_{hj} is the number of households in this enumeration district.

Overall, the ex ante selection probability $Prob(i)$ for each household i is given by multiplying the two partial probabilities. This probability may be shown as:

$$Prob(i) = \frac{m_{hj} m_h}{N_h} = \frac{1}{w_{Di}}.$$

The design weight (w_{Di}) is calculated by inverting this probability. For example, a household with a probability of selection equal to 0.001 has a preliminary design weight of *1,000=1/0.001*, which is much higher than that used for a household with a probability of selection equal to 0.009, which would be *111=1/0.009*.

This procedure ensures that every household that has the same probability of selection within a stratum on account of the sample design also has the same design weight. The design weights vary across the strata, due to the differing assumptions on the willingness of a household to participate, which determine the SSUs to be drawn, and the different size of the strata as a result of the number of households.

Finally, although the sampling frame was carefully prepared and cleaned before sampling, our sample still included some ineligible (see box 2) or duplicated observations (see also section 4.6.2.13), for example company addresses, addresses of care homes or secondary residences. We flagged all such cases detected during the fieldwork as ineligible or duplicated in our sample by setting the design weights equal to zero. As a result, the design weight total decreased from about 4.62 million to 4.57 million.

Table 12 shows some statistics of the obtained HFCS design weights across Austria's provinces. Carinthia and Vienna are the provinces with the lowest median weight, which is plausible, as households living in urban areas were oversampled during the fifth HFCS wave in Austria because of their low willingness to participate, which would have created a bias had they not been reweighted downward using the design weights. Very similar results are found in wave 4.

The value of a household's design weight can be interpreted as the number of households in the sampling frame that is represented by this household. For example, the median household in Vienna represents 572 households in the sampling frame.

Table 12

HFCS design weights by federal province

	Mean	Median	Minimum	Maximum
Vienna	563	572	0	618
Lower Austria	622	628	0	940
Burgenland	623	606	0	915
Styria	603	588	0	865
Carinthia	577	564	0	778
Upper Austria	629	617	0	870
Salzburg	624	581	0	943
Tyrol	598	602	0	776
Vorarlberg	584	615	0	801
Total	**598**	**581**	**0**	**943**

Source: HFCS Austria 2023, OeNB.

7.2.3 Nonresponse weights

As described in box 2, not all households participated successfully in the survey. If household characteristics correlate with nonresponse, the respondent population is not a random subsample of the sampling frame and the sample is nonresponse biased (see figure 5). In the HFCS, this is indeed the case, as can be seen in table 13. The table shows a logit regression of household participation in the survey (1 if the household participated, otherwise 0) on a set of variables that explain participation in the survey. The results show on the one hand that households living in municipalities with a higher share of employees in the primary sector or with few or no graffiti in the neighborhood or living in districts with higher crime rates have a lower probability of participating. On the other hand, households that live in areas with higher average age or in the countryside had an increased probability of participating. This suggests that nonresponse is not random.

This bias can be corrected by using nonresponse weights, i.e. by attaching a higher nonresponse weight to households with a low probability of responding than to households with a high probability of responding. To calculate the response probabilities and the corresponding nonresponse weights, the weighting class adjustment method is combined with the model-based adjustment method (see Biemer and Christ, 2008). The weighting classes are chosen optimally using the method described by Haziza and Beaumont (2007). The algorithm can be summarized in the following three steps:

- Step 1: The logit regression model shown in table 13 was used to estimate the probability of response for each household (assuming that the household was selected into the sample).
- Step 2: These households' response probabilities were grouped into eight classes. The number of classes and their resultant sizes are chosen optimally in line with Haziza and Beaumont (2007). To this end, a k-means algorithm is used to cluster households into a prespecified number of response classes with low variance and similar size. Next, class indicators are used as explanatory variables for the response propensity based on an ordinary least squares (OLS) regression from the logit regression

model estimated in step 1. Beginning with one class, the number of classes is increased in an iterative process until the corrected R^2 of this OLS regression exceeds 95%. This is the case for eight classes in the fifth wave of the HFCS in Austria.[5] Finally, the average response propensity for each class was calculated (unweighted total number of respondent households/unweighted total number of households).[6]

- Step 3: The nonresponse weight of a class is obtained by inverting the average response propensity of the respective class.

The advantage of this approach is that it stabilizes the nonresponse weights because the response propensities predicted by the regression model vary widely and can contain extreme values.[7] Information collected through interviewer surveys (see section 3.8), e.g. their level of education and experience, was found to correlate strongly and statistically significantly with households' response propensity and was therefore used in step 1. Additionally, sample design information and municipal or district-level information was used, which may also explain willingness to participate with statistical significance.

Box 2

Unit nonresponse in the HFCS in Austria

In the fifth wave of the HFCS in Austria, successful interviews were conducted with 2,849 households from the gross sample, which comprised 7,636 addresses. The remaining 4,787 addresses were classified either as unit nonresponse cases (4,701 households), ineligible addresses (85 addresses) or addresses of unknown eligibility (1 address).

The unit nonresponse cases are households as defined in the HFCS that were not interviewed successfully for several reasons. The most common reason was that households actively refused to take part in the survey, either by refusing to be interviewed, breaking off the interview or by failing to keep the interview appointment and being subsequently unavailable for contact. This applied to a total of 4,462 households. Another reason was that no contact at all could be established with 75 households. The remaining 164 nonrespondents specified other reasons, such as illness or language barriers; or they resulted from ex post exclusion of interviews due to a high number of missing or unreliable values.

In addition, 85 addresses were classified as ineligible because they were not part of the target population, as they were, for instance, addresses of companies, empty buildings or second homes of households that could be reached via their main residence address. Finally, the eligibility status of one address was impossible to ascertain, as the interviewers were unable to reach or find it. In accordance with how the eligibility statuses of the rest of the observed addresses in the sample were distributed, this address was randomly chosen to be eligible.

Expressed in percentages, the eligibility rate in the HFCS sample ultimately came to 98.9% and the nonresponse rate of the eligible households amounted to 62.3%. This means that successful interviews were conducted with 37.7% of the eligible households in the HFCS sample. This figure is comparable to the last wave but still significantly below the level of the previous waves (see the methodological report of the last wave for more information on the changes due to COVID-19). About 59% of the eligible households actively refused to take part in the survey.

[5] In wave four, seven classes were yielded as the optimal number of response classes by the algorithm.

[6] The average response propensity is unweighted (with respect to the design weights) for efficiency reasons. See Little and Vartivarian (2003) for more details.

[7] Another problem of the use of simple logit regression models, as highlighted by Iannacchione et al. (1991), is that such modeling does not ensure that the weighted sample marginal distributions conform to the population marginal distribution.

Table 13

Response propensity estimates based on a logit regression model

Covariates	Coefficients
Paradata on the interview, place of residence and neighborhood	
Household interview order	0.000285
	(0.000207)
Building characteristics (reference group: detached single-family house)	
Semi-detached single-family house	-0.239
	(0.150)
Single-family townhouse	0.322**
	(0.151)
Residential farm building	0.865***
	(0.199)
Apartment in a (high-rise) apartment building	0.472***
	(0.0710)
Student dormitory/rented room	1.379
	(1.429)
Other type of building	0.0221
	(0.259)
Building design characteristics (reference group: premium)	
Very good	-0.365**
	(0.153)
Medium	-0.455***
	(0.157)
Basic	0.130
	(0.189)
Very basic	-0.137
	(0.267)
Location characteristics (reference group: city center)	
Between the city center and suburbs	-0.0235
	(0.0778)
Suburbs and city outskirts	-0.000257
	(0.0885)
Countryside	0.222**
	(0.0990)
Graffiti in the neighborhood (reference group: many)	
Location – graffiti = 2, some	-0.570
	(0.476)
Location – graffiti = 3, few	-1.066**
	(0.466)
Location – graffiti = 4, none at all	-0.985**
	(0.464)
Interviewer characteristics	
Female interviewer	0.0821
	(0.0556)
Interviewer's age	0.00346
	(0.00255)
Interviewer's second-stage tertiary education	0.103
	(0.0637)
Interviewer's working experience in months	2.66e-05
	(0.000304)
Variables at the municipality level	
Average per capita income per municipality in 2020	-4.98e-06
	(8.95e-06)
Share of employees in the primary sector per municipality in 2019	-0.00824**
	(0.00416)
Share of university-trained population per municipality in 2019	-0.00209
	(0.00484)
Unemployment rate per municipality in 2019	-0.0145
	(0.00984)
Average age of population per municipality in 2021	0.0397***
	(0.0137)
Variables at the district level	
Average crime rate per district in 2021	-0.00457***
	(0.00128)
Constant	-0.938
	(0.786)
Observations[1]	7,551

Source: HFCS Austria 2023, OeNB

[1] The remaining 85 observations in the dataset are ineligible and are therefore not included in the regression.

Note: Standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1.

Table 14

HFCS nonresponse weights by response propensity

Response Classes	Predicted response propensity	Weight
	%	
I	0 to 25	3.798
II	25 to 31	3.442
III	31 to 35	3.193
IV	35 to 39	2.666
V	39 to 44	2.274
VI	44 to 49	2.174
VII	49 to 58	1.802
VIII	58 to 100	1.699

Source: HFCS Austria 2023, OeNB.

The HFCS nonresponse weights are shown in table 14. A value was calculated for each of the seven response groups and, by design, households with a high response propensity were assigned a lower weight than those with a low response propensity. Nonrespondent households were assigned a nonresponse weight equal to zero.

7.2.4 Poststratification weights

Erroneous exclusion may – as mentioned above – be an imperfection in the HFCS frame with respect to the target population. We may have missed households without postal addresses, which means that these types of households would be under-represented. If an external dataset covering these households and all others in our target population existed, we could use it to adapt our sample to this external dataset accordingly; we could put more weight on households without postal addresses so that the estimated size of the target population in the HFCS would be the same as the one in the external dataset.

Unfortunately, such a dataset does not exist in Austria. Similar surveys, like the EU SILC (EU Statistics on Income, and Living Conditions) or the Austrian microcensus, target a different population of households due to their specific household definition. While the target population of the HFCS includes all households (according to the above definition), the EU SILC and the Austrian microcensus only include households living in a dwelling

officially registered in the central residence registry as their main residence. This definition excludes a subset of households included in the HFCS household definition, namely all households living in dwellings that are not registered as a main residence or not registered at all. There are various reasons as to why in some cases households' actual main residences are not registered as such. For instance, students studying away from home may keep their main residence at their parents' address even though they are already a household of their own according to the HFCS definition; others may have just forgotten to register the address where they actually live as their main residence. Statistics Austria also acknowledges these problems and others when using main residence addresses for sampling households via the Austrian residence registry.[8]

Given that these datasets also suffer from erroneous exclusion, it does not make sense to reweight the *entire* sample according to the target population size of these datasets. However, following an adjustment made in the second wave of the HFCS in Austria, the survey establishes whether a household's residence is registered as a main residence or not, so it would appear to make sense to reweight this group of households to the microcensus. In particular, this may deliver a better picture with regard to the proportions of households in the Austrian provinces, as the Austrian microcensus uses a much larger sample than the HFCS. For the small group of remaining households in the HFCS sample that are not registered at their main residence, reweighting the microcensus does not seem sensible. Yet the erroneous exclusion bias in the HFCS sample is likely to be very small in this case, as the vast majority of households do have postal addresses. We constructed poststratification weights that put more weight on households with a lower probability of being included in the frame and less weight on households with a higher probability. We adjusted the HFCS sampling frame size only for households registered at their main residence according to the microcensus. Households not registered at their main residence are then added.[9] This increases comparability between the HFCS and the microcensus since the second wave and at the same time reduces the erroneous exclusion bias. Furthermore, poststratification weights can also reduce the sampling variance and, hence, increase the precision of the estimators; moreover, they can eliminate sample-specific random misrepresentations of the target population (see section 7.1).

The HFCS poststratification weights are constructed following the poststratification cell adjustment method (Biemer and Christ, 2008) and using the Austrian microcensus data (Q4 22) available during the field phase of the HFCS in Austria. The procedure was as follows:

- Step 1: Choose suitable predictors for including a household in the HFCS frame and cross-tabulating these variables to compute the poststratification cells. Different poststratification cells were defined

---

[8] For the microcensus, see Haslinger und Kytir (2006), p. 512 f; for the EU SILC, see Statistics Austria (2018), p. 44–45.

[9] Before the poststratification adjustment, the HFCS sampling frame encompassed 4,569,380 households, consisting of households registered at their main residence (4,541,504) and households not registered at their main residence (27,876 or about 0.6% of households). In wave 5, the registration status of two households could not be determined, as they responded with "don't know" or gave no answer to the question about being registered at their main residence. Based on the distribution of registration statuses among the observed households in the sample, both households were treated as registered at their main residence for poststratification adjustment purposes only. In the final dataset, this variable is going to be imputed like any other variable with missing values. After the poststratification adjustment of all households registered at their main residence, the population of these households comes to 4,067,474, which corresponds to the household population according to the 2022 Q4 microcensus. As a result, the final HFCS household population amounts to 4,095,350 (= 4,067,474 + 27,876).

depending on the registration status. For households registered at their main residence, the province, the tenure status of the main residence, household size and the age of the reference person serve as poststratification variables. No poststratification is performed for all other households, because, as described above, they are not included in the external dataset.

- Step 2: Calculate the average propensity to be included in the sampling frame for each cell.

$$\frac{HFCS\ frame\ population\ in\ the\ cell}{Microcensus\ frame\ population\ in\ the\ cell}$$

- Step 3: In each cell, the propensity was adjusted by a constant factor, thus adjusting the external dataset total.
- Step 4: Obtain the poststratification weight by inverting the average inclusion in-frame propensity for each cell.

Households registered at their main residence were grouped according to size, one group containing households with only one individual, another group containing those with two to four individuals and one containing those with five or more individuals.[10] This ensures that larger households are not underrepresented in the HFCS sample. Single households were further grouped according to age of the reference person, one group containing households with a reference person aged 0-64 and another group containing households with a reference person aged 65+. This ensures that older households are not overrepresented in the HFCS sample, which has become specially important since wave 4 (see the methodological report of wave 4 for more information on the changes due to COVID-19). Additionally, households registered at their main residence were grouped into tenants[11] and (partial) owners. Moreover, the households were assigned to the nine provinces.

Table 15 shows the HFCS poststratification weights – 67 values, i.e. one value per combination of registration status, province, tenure status of the main residence, household size and age of the reference person. The table shows, for example, that older households are overrepresented in the HFCS sampling frame, as they exhibit below-average poststratification weights.

7.2.5 Final weights

Three different weights were constructed to account for the different reasons as to why a household may misrepresent the target population. As we have seen, each of these weights can be interpreted as an inverted probability. The product of these yields a new inverted probability, which is the HFCS final weight w_i:

$$w_i = w_{Di} \cdot w_{NRi} \cdot w_{PSi}$$

$$w_i = \frac{1}{Prob(i\ is\ selected) \cdot Prob(i\ responds\,|\,i\ is\ selected) \cdot Prob(i\ is\ included\ in\ the\ frame)\,|\,i\ is\ selected\ and\ responds)}$$

$$w_i = \frac{1}{Prob(i\ is\ selected\ and\ i\ responds\ and\ i\ is\ included\ in\ the\ frame)}$$

The final weight w_i incorporates all three adjustments and can be interpreted as the inverted probability that household i is in the net sample. Households with a high probability of being in the net sample have

[10] Given the very low number in a poststratification cell, the cells were aggregated by household size in Vienna, Burgenland, Vorarlberg and Tyrol for main residence tenants, and in Burgenland and Vienna for main residence homeowners.

[11] Includes free users of main residences.

Table 15

HFCS poststratification weights for registration status of main residence, province, tenure status, household size and age by response propensity

	Official main residence								Other households							
	Household size (number of persons)								Household size (number of persons)							
	1		2 to 4	5 or more	1		2 to 4¹	5 or more¹	1		2 to 4	5 or more	1		2 to 4	5 or more
	Age of household's reference person				Age of household's reference person				Age of household's reference person				Age of household's reference person			
	0-64	65+			0-64	65+			0-64	65+			0-64	65+		
	Homeowners				Tenants				Homeowners				Tenants			
Vienna	1.137	0.418	1.434		0.814	0.338	1.180	1.420								
Lower Austria	1.573	0.725	0.886	4.829	0.791	0.473	0.742	5.203								
Burgenland	3.233		0.551		1.617	0.644	2.303									
Styria	1.649	1.435	1.352	2.853	0.655	0.592	0.597	3.326								
Carinthia	2.359	0.603	1.283	1.468	0.728	0.683	0.807	2.069	1							
Upper Austria	0.901	0.724	1.055	2.357	0.967	0.700	0.916	1.775								
Salzburg	1.130	0.819	1.221	1.930	1.237	0.444	0.629	1.815								
Tyrol	1.889	0.547	1.110	1.285	0.651	0.598	1.515									
Vorarlberg	0.751	0.378	1.137	2.235	0.544	0.715	1.504									

Source: HFCS Austria 2023, OeNB.

[1] Given the very low number in the poststratification cell "5 or more" for tenants in Vienna, we used the household size categories "2 to 3" and "4 or more" for these cases.

a lower final weight and represent fewer households in the target population than households with a low probability of being in the net sample.

The combination of nonresponse weights and poststratification weights results in 536 different weight adjustment cells based on registration status, province, tenure status of the main residence, household size, age of the reference person and the response propensity classes described above. Each household is represented in precisely one of these cells.

Finally, once we have taken the design weights into account, we obtain the HFCS final weights, whose distribution is shown in chart 1. The HFCS final weights range from 319 to 8,561, with the mean being 1,437 and the median 1,268. Their distribution is slightly skewed to the right, which is not atypical for unequal probability sample designs. After all, households with a higher probability of selection (below average design weights) dominate the sample. This effect is reinforced by the further weight adjustments.
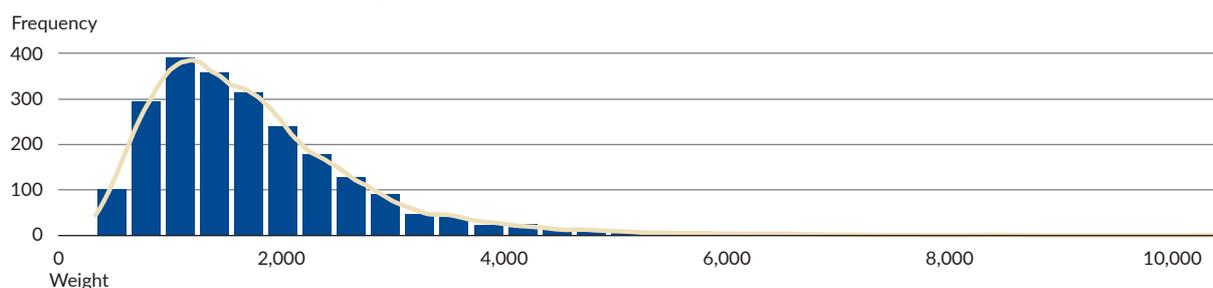
7.3 Selected results

Table 16 shows the impact of the HFCS final weights on estimations by comparing selected weighted and unweighted mean values of HFCS variables. For example, we can see that households in Vienna were slightly downweighted from 23.6% to 23.1%. This means that after their oversampling, overall households in Vienna were overrepresented in the sample with respect to the target population. The comparison also shows that households with higher income or higher net wealth were overrepresented in the unweighted sample, which is probably caused by these households' higher nonresponse rate.

The use of the final HFCS weights is sufficient when calculating the weighted statistics shown in table 16. To calculate the appropriate correct variances or standard errors of these estimators, however, replicate weights, which are described in chapter 8, are necessary.

Chart 1

Distribution of HFCS final weights

Source: HFCS Austria 2023, OeNB.

7.4 Concluding remarks

Table 16

Comparison of weighted and unweighted means of selected HFCS variables (imputed)

	Mean	
	Unweighted	Weighted
Household size (number of persons)	1.86	2.09
	% of households	
Vienna	23.6	23.1
Lower Austria	19.2	18.5
Burgenland	3.5	3.2
Styria	14.1	13.8
Carinthia	6.7	6.4
Upper Austria	14.8	16.2
Salzburg	5.9	6.1
Tyrol	8.0	8.5
Vorarlberg	4.3	4.3
	EUR	
Estimated household monthly net income	2,809	3,100
Household net wealth	275,506	336,633

Source: HFCS Austria 2023, OeNB.

We constructed a set of final weights to correct imperfections in the unweighted HFCS sample with respect to the HFCS target population. These imperfections are unequal probability sampling bias, erroneous inclusion, frame multiplicity and erroneous exclusion and nonresponse bias.

While the weighted HFCS sample enables unbiased population estimates, it also increases the variance of the population estimates, which makes them less precise.[12] According to the unequal weighting effect (UWE) statistic developed by Kish (1995), the variance of HFCS population estimates may be increased by a maximum of 26.1% *(UWE = 1 + (coefficient of variation)² = 1.261)* as a result of weighting. This figure is clearly below the level recorded in the last wave, yet it remains significantly above the levels observed in earlier waves (see the methodological report of the last wave for details on changes related to COVID-19). However, it is still not necessary to apply weight trimming methods. Furthermore, a small increase in variance is acceptable in return for a significant reduction in the bias if it helps to avoid distorted results being classified as significant too often.

An explanation of how to correctly use the weights in Stata and R is provided in chapter 9, User guide.

---

[12] The poststratification step can restrict this increase in sample variance (see Levy and Lemeshow, 2008).

8 Construction of replicate weights for variance estimation

8.1 Introduction

The use of the final survey weights described in chapter 7 is sufficient when estimating population parameters. However, to calculate the corresponding correct variances or standard errors of these estimators requires replicate weights, which are described here. HFCS sampling involves a variety of complex features, such as stratification, multistage sampling, proportional-to-size sampling in the first stage or sampling without replacement in the second stage. In addition, the design weights are adjusted for nonresponse and poststratification. Ignoring these features in statistical analysis will bias the estimated variances of point estimators. For example, if stratification is ignored, the standard errors will be too large, and if clusters are ignored, the standard errors will be too small. Furthermore, if design weights are ignored, the sampling distributions of the statistics underrepresent the observations with a low selection probability and over-represent those with a high selection probability (see Kolenikov, 2010).

A problem that occurs frequently when statistical analysis takes into account a complex survey design with all its features is that the mathematical functions of the variance estimators are unknown. Therefore, performing a statistical analysis requires methods developed especially for the purpose of variance estimation. There are two general categories of variance estimation methods: *replicate weight methods (also called replication or resampling methods) and linearization.*[1]

Until recently, linearization was preferred to replication in the literature, as linearization requires less computational power. However, linearization comes with the major disadvantage that data protection regulations prevent some information necessary for linearization from being shared. When replicate weights are used, data availability and privacy are not an issue. After all, replicate weights consist of numerous variables which are not available to data users (e.g. stratum and primary sampling unit (PSU) variables). Leaving out this information makes it virtually impossible for individual respondents to be identified from the data by the data user (see Heeringa et al., 2017).

Moreover, the linearization method is unsuitable for estimating the variance of nonlinear statistics (medians, quartiles, etc.), as it requires computing derivatives of continuous functions; however, quantile functions, for instance, are discontinuous. Replicate weights, by contrast, are well suited for estimating the variance of such statistics (see Heeringa et al., 2017).

Given the data protection requirements mentioned above and because the HFCS data facilitate in particular the analysis of distributional parameters such as medians and quantiles, we use – in accordance with ECB requirements – replicate weights for variance estimation in the HFCS.[2] In the following section, we describe how replicate weights were constructed for the HFCS in Austria.

---

[1] For a comprehensive overview of variance estimation methods, see Levy and Lemeshow (2008) or Heeringa et al. (2017).

[2] In combination with multiple imputations, variance estimation of nonlinear statistics by means of resampling-weights is still largely unexplored.

8.2 Construction of replicate weights

8.2.1 The replication method

The replication method aims to estimate the variance of an estimated population parameter. The idea behind this is, to begin with, to estimate population parameters for individual subsets (so-called replicates) of the sample observations. The variability of these estimated population parameters across all replicates is subsequently calculated, resulting in the desired variance of the estimated population parameter (see Levy and Lemeshow, 2008).

Instead of saving a whole sample for each replicate, it is more practical to vary the final survey weights. For example, instead of removing a sample observation to construct a certain replicate, it can be given a weight of zero in the replicate. Then the weights of the other observations in the same stratum need to be increased to ensure that the totals are unbiased for each replicate r (see Kolenikov, 2010). The replicate weights $w_i^{(r)}$ for $r = 1,..., R$ will be published together with the HFCS dataset.

There are different methods to form such replicates. The three major replication methods used in survey literature are *balanced repeated replication, jackknife repeated replication and bootstrap replication.* Although in most simulations in the literature, the estimators of all three replication methods converge towards one another as the sample size increases, simulation studies have shown that bootstrap and balanced repeated replication are better suited to quantile estimation than jackknife (see Kovar et al., 1988). Finally, as balanced repeated replication works only in designs with exactly two PSUs per stratum, which is not the case in the HFCS in Austria, we decided to use the *(rescaling) bootstrap procedure* proposed by Rao and Wu (1988) and enhanced by Rao et al. (1992). This procedure is also in line with the provisions of the ECB's Household Finance and Consumption Network.

The bootstrap procedure forms replicates based on repeated with-replacement sampling of the PSUs within a stratum. The idea is to mimic the original sampling procedure in order to obtain approximations for the sampling distributions of the relevant statistics.

8.2.2 Sampling error calculation model

To mimic the original sampling procedure, we create a sampling error calculation model that is a simplification (see Heeringa et al., 2017) of the actual complex sample design (see chapter 6).

In the HFCS in Austria, one necessary simplification of the sampling error calculation model compared with the original sampling procedure is to collapse, i.e. merge, strata with one single PSU because the bootstrap procedure requires at least two PSUs per stratum. Due to the specific stratification of the HFCS sample design, single-PSU strata are quite common in the sample: Only one PSU was drawn from 64 out of 192 strata. For the sampling error calculation model, every single-PSU stratum is paired with the geographically nearest stratum to form a single pseudo stratum, taking into account how many PSUs are in this stratum. Aggregation is carried out with the nearest stratum containing a smaller number of PSUs, reducing the frequency of necessary aggregations. Although collapsing the strata produces an upward bias in the estimated variance, this bias is kept as small as possible by collapsing geographically close strata, which keeps the PSUs within one pseudo-stratum very homogeneous. In this context it must be pointed out that upward biases of standard errors lead to a loss in statistical power. In general, however, this is more acceptable than negative biases of standard errors, which lead to results that are too often considered statistically significant.

Table 17 shows how stratum size (number of PSUs drawn per stratum) changes when the HFCS sampling error calculation model is used instead of the original HFCS sample design: When collapsing strata in the sampling error calculation model, their number decreases from 192 to 151, which means stratification is still very high. Moreover, the mean stratum size increases from 3.6 PSUs to 4.6 PSUs per stratum.

Comparison of HFCS design strata and HFCS pseudo strata

	Design strata	Pseudo strata
Number of strata	192	151
Mean size	3.6	4.6
Median size	2	3
Minimum size	1	2
Maximum size	36	36

Source: HFCS Austria 2023, OeNB.

Note: Stratum size as measured by PSUs drawn per stratum.

Another simplification performed in the HFCS sampling error calculation model in contrast to the original sample design is to assume that sampling variance stems mostly from the first stage of sampling (i.e. the selection of PSUs, and not that of households within each PSU). Therefore, two-stage sampling is reduced to single-stage sampling where all gross sample households within drawn PSUs are selected in the replicate sample.

In addition, all PSUs have the same probability of being selected into the replicate sample. Thus, the sampling error calculation model simplifies sampling by making a PSU's probability of being drawn independent of its size as measured by the number of households.

No further simplifications are required by the sampling error calculation model. The nonresponse and poststratification weight adjustments are implemented in the same way as in the original weighting procedures (see chapter 7), and a finite population correction[3] is performed.

8.2.3 Construction of replicate weights

The algorithm used to construct the HFCS replicate weights comprises the following steps:
- Step 1: Draw m_h PSUs with replacement within each pseudo stratum h.
- Step 2: Adjust the final survey weights of the drawn observations to create a new set of replicate weights. In particular, apply the same nonresponse and poststratification weight adjustments (sections 7.2.3 and 7.2.4) as for the final design weights and perform a finite population correction.
- Step 3: Repeat steps 1 and 2 R times to obtain $r = 1,..., R$ sets of replicate weights.

In step 1, the number of PSUs m_h drawn in each stratum of size n_h is set to $m_h = n_h - 1$. This decision is taken often to ensure the efficiency of the bootstrap estimators without violating the natural parameter ranges (see Kolenikov, 2010).

In step 2, the final survey weights must be adjusted because some PSUs may be duplicates and some may not have been drawn at all. As a consequence, each replicate will be biased with respect to the target population and therefore, to obtain the replicate weights, the design weights must be adjusted in the same way they were adjusted when constructing the final survey weights (see chapter 7). In addition, a finite population correction is required, as PSUs are sampled without replacement in the original HFCS sample design (see footnote 3).[4]

---

[3] The finite population correction accounts for the reduction in variance that occurs when sampling without replacement from a finite population. This type of sampling is used in the sample design of the second stage of the HFCS in Austria.

[4] In the HFCS sample design, PSUs are drawn with replacement, SSUs without. Although the sampling error calculation model ignores the second stage, a finite population correction was performed to allow for the fact that households are not allowed to appear twice in the sample. Finite population correction reduces the bias of a higher variability of replicate weights.

Finally, in step 3, the higher the number of replicates R is, the more precise the standard error estimates are. We choose $R = 1,000$, which lies in the upper bound of the usual recommendations found in literature (see Kolenikov, 2010).

Table 18 shows some descriptive statistics of a selection of HFCS replicate weights. We can see that owing to the homogeneous weighting adjustments, the mean and the total sum of replicate weights remain unchanged. Moreover, compared with the final survey weights in the HFCS, the replicate weights have smaller minimum values, however none are equal to zero. These values correspond to the nonselected PSUs, which, instead of being assigned a weight equal to zero, are assigned a small positive weight in the finite population correction. The fact that the replicate weights also have larger maximum values than the final survey weights can be explained by the weight adjustments that were carried out: As some PSUs are not drawn in the replicates, and in order to obtain the same estimated population sizes as in the original sample, the weights of the observations in the drawn PSUs must be increased.

Table 18

Selected HFCS replicate weights

	Mean	Median	Minimum	Maximum	Total
Final Survey weights	1,437	1,268	320	8,562	4,095,350
1st set of replicate weights	1,437	1,062	7	18,850	4,095,350
2nd set of replicate weights	1,437	1,088	8	23,689	4,095,350
3rd set of replicate weights	1,437	1,156	7	13,295	4,095,350
998th set of replicate weights	1,437	1,109	7	25,387	4,095,350
999th set of replicate weights	1,437	1,183	9	18,196	4,095,350
1,000th set of replicate weights	1,437	853	8	22,771	4,095,350

Source: HFCS Austria 2023, OeNB.

Note: Statistics refer to successfully interviewed households only.

8.3 Concluding remarks

We constructed 1,000 sets of replicate weights to enable HFCS data users to correctly estimate the standard errors of point estimators in the HFCS. This is necessary because the complex features of the HFCS survey design, which comprises amongst other things stratification, several stages of cluster sampling and weighting adjustments, bias the variance estimators if data users ignore them.

While it is true that correctly calculating the standard errors by using replicate weights requires more computational power than analyzing the data without using replicate weights, in practice it is not necessary to use all 1,000 sets of replicate weights for variance estimation. Thus, for example, it is possible to perform variance estimations using fewer replicates more quickly but less precisely. The number of replicates used depends on the type of estimator and the size of the population surveyed (see e.g. Pattengale et al., 2010). For instance, estimating the means for the total population will, as a rule, require fewer replicates than estimating the medians for specific population subgroups.

See the HFCS User guide (chapter 9) for an explanation of how to use the replicate weights correctly in Stata and R.

9 User guide

9.1 Introduction

As we have seen in the previous chapters, the HFCS data are characterized by special features that must be taken into account when analyzing the data. The data are multiply imputed and contain survey weights and replicate weights. The HFCS data are also stored in several files, due to the structure of the survey. These files differ in terms of the data level (household or individual), the number of implicates (i.e. each implicate is a separate file) and the type of data, depending on whether the data were collected or constructed (derived variables, i.e. aggregated variables, and replicate weights vs. survey variables). This chapter[1] provides Stata[2] code that users can employ step by step to account for all of these features.[3] The ECB is expected to also make several program codes available in summer 2026 when publishing the dataset. In this chapter Stata program code is contained in the blue boxes. When copied into the Stata command window[4], it must be run in the sequence outlined below; altering the sequence may corrupt the code. Additionally, the online appendix contains a do-file "`user_guide.do`" with all the steps that are laid out below.[5] We also include an R version of the user guide (called "`user_guide_r_1.0.R`") in the online appendix, which provides a code for the R environment that is very similar to what is explained for Stata below. Since the steps are similar, we refrain from a repetition of detailed explanations in the chapter. This chapter first explains how to merge the separate files, then describes one way to set up the structure for imputations and survey information. Finally, some examples of simple estimation commands and how they are used are provided.

---

[1] The authors refrain from making a judgment about which programs to use and with which settings. In particular if the size of the subsamples varies in each iteration, the estimation of discontinuous estimators does not comply with the assumptions of the results evidenced in literature (e.g. Little and Rubin, 2002). It is the responsibility of users to check whether individual estimation commands are valid and adequate under particular conditions.

[2] The codes were written for Stata version 15.1 or higher and may not be valid for previous Stata versions. The code extracts were provided by Sébastien Pérez-Duarte (ECB) and have been slightly altered and expanded for release here.

[3] Any changes and improvements made to the code are continuously updated in the online appendix. Any adjustments made since the release of the first wave of the HFCS were included in this program code.

[4] Due to the way Stata handles line breaks, they may need to be deleted if the program code is copied by hand.

[5] The two macros containing the individual path to the data and the additional do-files must be specified before execution. Given the size and structure of the data and depending on software and hardware specifications, executing the do-file may require a long time.

9.2 Merging the data files

The core HFCS data, which contain all internationally agreed variables, consist of the five multiply imputed samples or implicates at the household level (files H1–H5), the corresponding samples at the individual level (files P1–P5) and the corresponding set of aggregated variables[6] (files D1–D5). Before creating a new dataset containing all these files, users must specify the path to the datasets and the folder containing the do-files on their computers. The variables used for merging are the household identifier "sa0010," the implicate number "im0100" and the country identifier "sa0100."

```
************************************************************************
***Merging the files of the HFCS data
************************************************************************
*Set macro for the path to the data (must be specified by the user)
global hfcsdata="path to the appropriate folder where the data are stored"

*Set macro for the path to the do-files (must be specified by the user)
global hfcsdofile="path to the appropriate folder where the do-files are stored"

*Set working directory
cd "$hfcsdata"

*Reshaping and merging the p and h files together (wide format)
forvalues i=1(1)5 {
  use "$hfcsdata\P`i'.dta", clear
  drop id hid survey
  foreach var of varlist sa0010- fra0500 {
    local `var'lab: variable label `var'
  }
  gen idpers_temp1="_"
  egen idpers_temp2=concat(idpers_temp1 ra0010)
  drop ra0010
  reshape wide ra0?0* fra0?0* ra0020 fra0020 ra0030 fra0030 ra0040 fra0040 p* fp* ///
  , i(sa0010 sa0100) j(idpers_temp2) string
  drop idpers_temp*
  foreach j of varlist ra* fra* p* fp* {
    local last2car=substr("`j'", `=length("`j'")-1', 1)
    local last1car=substr("`j'", length("`j'"), 1)
    if "`last2car'"=="1" {
      local firstcar=substr("`j'",1, `=length("`j'")-3')
      label variable `firstcar'_`last2car'`last1car' ///
      "``firstcar'lab' - `last2car'`last1car'"
    }
    else {
      local firstcar=substr("`j'",1, `=length("`j'")-2')
      label variable `firstcar'_`last1car' "``firstcar'lab' - `last1car'"
    }
  }
}
```

---

[6] The ECB is expected to make the definitions of the aggregated variables and the datasets available in summer 2026.

```
save "$hfcsdata\P`i'_temp.dta", replace
  clear
  use "$hfcsdata\H`i'.dta", clear
  merge 1:1 sa0010 sa0100 im0100 using "$hfcsdata\P`i'_temp.dta", nogen
  save "$hfcsdata\M`i'.dta", replace
  erase "$hfcsdata\P`i'_temp.dta"
}

*Merging the core with the derived variables
forvalues i=1(1)5 {

  use "$hfcsdata\M`i'.dta", clear
  merge 1:1 sa0010 im0100 sa0100 using "$hfcsdata\D`i'.dta"
  save "$hfcsdata\temp`i'.dta", replace
}

*Merging the implicates together1
use "$hfcsdata\temp1.dta", clear
forvalues j=2(1)5 {
  append using "$hfcsdata\temp`j'.dta"
}

*Drop unnecessary variables and labels
drop _merge
label drop _merge

*Save the HFCS data
save "$hfcsdata\hfcs.dta", replace
```

1 The temporary files are kept for configuring the multiple imputed data and are only erased following this procedure.

So-called M-files are created by reshaping the P-files (using the `reshape` command), including the appropriate labeling of the P-file variables, and by merging the resulting dataset with the H-files. The reshaping of the P-Files is based on a temporary string variable for a correct naming of the personal level variables. The M-files are provided in *wide format*,[7] i.e. one line of the data matrix contains information on a given household and the information on each individual within a household is included in a separate variable. Merged with the D-files, these M-files yield the entire HFCS dataset in the "hfcs.dta" file.

---

[7] It is also possible to merge the data files in "long" format using an almost identical code without needing to reshape the personal files.

9.3 Multiple imputations

The next step is to import both the original data and the imputed values into Stata's `mi` (i.e. `mi esti-mate` commands for appropriate use of the multiple imputation structure). As the original data are not part of the HFCS data files, we have to construct them from the information about whether observations vary across implicates (indicating multiple imputation and, hence, missing values) and from the information about missing values taken from the flags.[8] Finally, original and imputed data must be imported and registered. Users should take note of the "IMPUTEDVARS" macro in the program code below, which contains a string listing all imputed variables once the corresponding loop has been executed. Moreover, the aggregated variables are registered as having been passively imputed. If registration was successful, running the `mi varying` command should yield only a few variables (e.g. the implicate number "im0100") and the flags should be shown as "unregistered varying."

```
************************************************************************
***Preparing the data for mi import
************************************************************************
*Create the zero implicate to simulate the original data
*Use one implicate of the data
use "$hfcsdata\temp1.dta", clear
*Replace the implicate number by "0" to simulate the original data
replace im0100=0
*Append all other implicates
append using "$hfcsdata\hfcs.dta"

*For some reason string variables do not play well with mi commands and need to be encoded ///
into numeric variables
foreach var of varlist hb* hc* hd* hg* hh* hi* /*hr* */ pa* pe* pf* pg* ra* sa0100 sb1000 {
  capture confirm numeric variable `var'
  if _rc {
    rename `var' `var'_string
    encode `var'_string, gen(`var')
    drop `var'_string
  }
}

*Set as soft missing (".") in im0100==0 all values varying, and also those whose flags set ///
them as imputed
global IMPUTEDVARS=""
```

[8] All missing values (including "Don't know," "No answer" and skip patterns) are set to "." and are paired with specific flags reflecting different types of missing values (e.g. skipped observations are flagged with a "0"). Flag variables have the same variable name, but their names are preceded by an "f."

```
foreach var of varlist hb* hc* hd* hg* hh* hi* /*hr* */ pa* pe* pf* pg* ra* {
  capture confirm numeric variable `var'
  if !_rc {
    tempvar sd count
    quietly bysort sa0100 sa0010 : egen `sd'=sd(`var')
    quietly bysort sa0100 sa0010 : egen `count'=count(`var')
    quietly count if ( (`sd'>0 & `sd' <. ) | `count'<6 | ///
    (f`var'>4000 & f`var'<5000) ) & im0100==0
    if r(N)>0 global IMPUTEDVARS "$IMPUTEDVARS `var'"
    quietly replace `var'=. if ( (`sd'>0 & `sd' <. ) | `count'<6 | ///
    (f`var'>4000 & f`var'<5000) ) & im0100==0
    drop `sd' `count'
    disp ".", _continue
  }
}

*Here we need to set all derived variables for im0100==0 missing because it is passively imputed
foreach var of varlist d* hb3001-hb40033 hb4099 hb4105 hb4205 {
  local type1: type `var'
  local type2=substr("`type1'",1,3)
  if "`type2'"!="str" {
    replace `var'=. if im0100==0
  }
}


*Drop unnecessary variables
drop id _merge

*Save the HFCS data
save "$hfcsdata\hfcs.dta", replace

*Erase temporary files that will not be needed anymore
forvalues i=1(1)5 {
  erase "$hfcsdata\temp`i'.dta"
}

**********************************************************************
*****Import as multiply imputed data
**********************************************************************
*Import the imputation structure of the data into Stata
mi import flong, m(im0100) id(sa0100 sa0010) clear

*Register the variables that are imputed
mi register imputed $IMPUTEDVARS

*Register derived variables as passively imputed
mi register passive d*

*Check whether all imputed variables are registered
mi varying

*Save the HFCS-data with mi structure
save "$hfcsdata\hfcs.dta", replace
```

9.4 Survey variables

Having configured the data as multiply imputed, we can designate the data as complex survey data, identify variables that contain information about the survey design and specify the default method for variance estimation. In our case, all this information is contained in the final survey weights (hw0010) and in the 1,000 sets of replicate weights (wr0001–wr1000), which are provided in a separate file and hence have to be merged with the data first.

```
**********************************************************************
***Setting up Complex Survey Design
**********************************************************************
*Encode country indicator
use "$hfcsdata\W.dta", clear
rename sa0100 sa0100_string
encode sa0100_string, gen(sa0100)
drop sa0100_string
save "$hfcsdata\Wtemp.dta", replace

*Using the HFCS data with mi structure
use "$hfcsdata\hfcs.dta", clear
*Merging the data with replicate weights
merge m:1 sa0100 sa0010 using "$hfcsdata\Wtemp.dta"

*Drop unnecessary variable and files
drop _merge
erase "$hfcsdata\Wtemp.dta"

*Setting the appropriate survey structure using replicate weights
mi svyset [pw=hw0010], bsrweight(wr0001-wr1000) vce(bootstrap)

*Save the HFCS-data with mi svyset structure
save "$hfcsdata\hfcs.dta", replace
```

9.5 Standard estimation procedures

The data are now ready to be analyzed in Stata. After writing "`mi estimate: svy:`" followed by the estimation command in question, Stata will provide correct estimates and standard errors, taking into account both the multiple imputation framework and the replicate weights.[9] The `esampvaryok` option can be useful when the sample size varies across implicates due to imputations.[10] Stata versions below Stata 12 do not allow the use of replicate weights together with multiply imputed data. For versions above Stata 12, the option `vceok` (used after the `mi estimate` command, e.g. "`mi estimate, vceok:…`") can be used as a workaround. It should be noted that in order to calculate the correct variance for subsamples of households (see second example in the following program code), Stata requires a dummy variable for each of these subsamples combined with the use of the option for subpopulations ("`…svy, subpop(dummy)…`").[11] Alternatively, it is possible to use the option over(variable) for certain estimation commands (see last example in the following program code).

```
**********************************************************************
***Using Standard Estimation Procedures
**********************************************************************
*Using the HFCS-data with mi svyset structure
use "$hfcsdata\hfcs.dta", clear

*Mean of current value of primary housing unit
mi estimate, esampvaryok vceok: svy: mean hb0900

*Mean of current value of primary housing unit for part owner of the primary housing unit
gen partowner=(hb0300==2)
mi estimate, esampvaryok vceok: svy, subpop(partowner): mean hb0900

*Proportions of owner/renter of primary housing unit
mi estimate, esampvaryok vceok: svy: proportion hb0300

*Ratio of current to acquisition value of primary housing unit
mi estimate, esampvaryok vceok: svy: ratio hb0900 hb0800

*Regression of current value of primary housing on acquisition value and year of acquisition
mi estimate, esampvaryok vceok: svy: regress hb0900 hb0800 hb0700

*Average level deposits according to gender of the first person
mi estimate, esampvaryok vceok: svy: mean da2101, over(ra0200_1)
```

[9] A correct point estimate of statistics can be carried out on the basis of the final survey weights. Replicate weights are needed to calculate a variance estimator.

[10] Rubin's combination rules (see e.g. Little and Rubin, 2019) were derived on the assumption that the same set of observations is used in each imputed data set. Thus, they may not necessarily apply when the sets of observations used in the data analysis differ. This is why `mi estimate` generates an error when this happens. When the subsets used in each complete data analysis differ relatively little, the conventional formulas may still be applicable. In this case, users can choose to use the esampvaryok option or find a better way to deal with the violation of the assumption of Rubin's combination rules described above. To our knowledge, this issue has not yet been addressed in literature.

[11] The use of an if-condition does not account for the uncertainty of the subsample size and therefore yields incorrect variance estimators.

9.6 Additional estimation procedures

To calculate medians or other quantiles, we use a different Stata package, called `medianize`, which was developed by the ECB (the respective do-file can be found in the online appendix). It must be used with caution as it is not yet a standard feature of Stata; so far it has been tested only in limited environments. Other Stata features used are the tabstat command and analytical weights.

```
**********************************************************************
***Including Additional Estimation Procedures
**********************************************************************
*ECB-written command to calculate medians (and some other quantile statistics), which should be ///
run before the estimation command
capture program drop medianize
do "$hfcsdofile\medianize.do"

*Median of amount still owned in the first loan collateralized with primary housing unit
mi estimate, esampvaryok vceok: svy: medianize hb1701

*Median of amount still owned in the first loan collateralized with primary housing unit over ///
gender of first person
mi estimate, esampvaryok vceok: svy: medianize hb1701, over(ra0200_1)

*Median of amount still owned in the first loan collateralized with primary housing unit over ///
gender of first person
mi estimate, esampvaryok vceok: svy: medianize hb1701, over(ra0200_1) stat(p10)
```

9.7 Online appendix

The online appendix contains the Stata code described above and the do-files necessary to estimate certain quantiles. Additionally, an R version is provided. The code in the online appendix will be updated as required to include the user guide program codes for other HFCS-relevant topics.

References

Albacete, N. 2014. Multiple imputation in the Austrian Household Survey on Housing Wealth. In: Austrian Journal of Statistics. 43(1). 5–28.

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. OeNB. 57–73.

Albacete, N. and P. Lindner. 2015. Foreign currency borrowers in Austria – evidence from the Household Finance and Consumption Survey. In: Financial Stability Report 29. OeNB. 93–109.

Albacete, N. and P. Lindner. 2017a. Simulating the impact of borrower-based macroprudential policies on mortgages and the real estate sector in Austria – evidence from the Household Finance and Consumption Survey 2014. In: Financial Stability Report 33. OeNB. 52–68.

Albacete, N. and P. Lindner. 2017b. How strong is the wealth channel of monetary policy transmission? A microeconometric evaluation for Austria. In: Monetary Policy & the Economy Q2/17. OeNB. 32–53.

Albacete, N. and M. Schürz. 2013a. Vergleich der Einkommensmessung für Haushalte in Österreich: HFCS versus EU-SILC. In: Statistiken – Daten und Analysen Q2/13. OeNB. 88–89.

Albacete, N. and M. Schürz. 2013b. Interviewereffekte beim HFCS Austria 2010. In: Statistiken – Daten und Analysen Q3/13. OeNB. 57–68.

Albacete, N. and M. Schürz. 2014a. Paradaten im HFCS Austria 2010 – Teil 1: Evaluierung von Non-Response-Fehlern. In: Statistiken – Daten und Analysen Q1/14. OeNB. 81–97.

Albacete, N. and M. Schürz. 2014b. Paradaten im HFCS Austria 2010 – Teil 2: Evaluierung von Messfehlern. In: Statistiken – Daten und Analysen Q3/14. OeNB. 54–64.

Albacete, N. and M. Schürz. 2015. Interviewereffekte auf Haushaltsvermögen am Beispiel des Household Finance and Consumption Survey Austria 2010. In: Statistiken – Daten und Analysen Q4/15. OeNB. 55–63.

Albacete, N., M. Andreasch and P. Lindner. 2018. Verschuldung der privaten Haushalte in Österreich. In: Statistiken Sonderheft. OeNB. June.

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

Albacete, N., P. Fessler and P. Lindner. 2016. The distribution of residential property price changes across homeowners and its implications for financial stability in Austria. In: Financial Stability Report 31. OeNB. 62–81.

Albacete, N., P. Fessler and P. Lindner. 2018. One policy to rule them all? On the effectiveness of LTV, DTI and DSTI ratio limits as macroprudential policy tools. In: Financial Stability Report 35. OeNB. 67–83.

Albacete, N., P. Fessler and P. Lindner. 2025a. The Wealth Distribution and Redistributive Preferences: Evidence from a Randomized Survey Experiment. In: European Journal of Political Economy, Volume 90, Part B. 102741. Forthcoming.

Albacete, N., P. Fessler and P. Lindner. 2025b. Who's asking? Interviewer effects on unit nonresponse in the Household Finance and Consumption Survey. In: Survey Methodology. Forthcoming.

Albacete, N., P. Fessler and M. Propst. 2020. Mapping financial vulnerability in CESEE: understanding risk-bearing capacities of households is key in times of crisis. In: Financial Stability Report 39. July. https://www.oenb.at/dam/jcr:fe93c400-de9f-4163-9644-66b54afc8d60/09_Mapping_financial_vulnerability_in_CESEE.pdf (accessed on October 2, 2025).

Albacete, N., S. T. Dippenaar, P. Lindner and K. Wagner. 2019. Eurosystem Household Finance and Consumption Survey 2017. Methodological notes for Austria. Monetary Policy & the Economy Q4/18 – Addendum. OeNB.

Albacete, N., P. Fessler, F. Kalleitner and P. Lindner. 2021. How has COVID-19 affected the financial situation of households in Austria? In: Monetary Policy and the Economy Q4/20 – Q1/21. OeNB. 111–130.

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

Albacete, N., J. Eidenberger, G. Krenn, P. Lindner and M. Sigmund. 2014. Risk-Bearing Capacity of Households – Linking Micro-Level Data to the Macroprudential Toolkit. In: Financial Stability Report 27. OeNB. 95–110.

Albacete, N., I. Gerstner, N. Geyer, P. Lindner, N. Prinz and V. Woharcik. 2022. Effects of interest rate and inflation shocks on household vulnerability in Austria: a microsimulation using HFCS data. In: Financial Stability Report 44. OeNB. November. 69–77.

Andreasch M., P. Fessler and M. Schürz. 2013. HFCS des Eurosystems – Möglichkeiten und Einschränkungen von Ländervergleichen im Euroraum. In: Statistische Nachrichten 9/2013. 842–51.

Banca d'Italia. 2012. Sample surveys – Household income and wealth in 2010. Supplements to the Statistical Bulletin XXII(6).

Barceló, C. 2006. Imputation of the 2002 wave of the Spanish Survey of Household Finances (EFF). Banco de España Documentos ocasionales 0603.

Beer, C. and K. Wagner. 2017. Households' housing expenditure in Austria, Germany and Italy. In: Monetary Policy & the Economy Q4/17. OeNB. 48–61.

Bekhtiar, K., P. Fessler and P. Lindner. 2019. Risky assets in Europe and the US: risk vulnerability, risk aversion and economic environment. ECB Working Paper Series No 2270. https://www.ecb.europa.eu/pub/pdf/scpwps/ecb.wp2270~9c72a27c18.en.pdf (accessed on October 2, 2025).

Biemer, P. and S. Christ. 2008. Constructing the survey weights. In: P. Levy and S. Lemeshow. Sampling of populations: Methods and applications. 4th edition. Wiley. 489–516.

Bledsoe, R. and G. Friess. 2002. Editing the 2001 Survey of Consumer Finances. Annual Meeting of the American Statistical Association. Joint Statistical Meetings. New York. August. 11–15.

Bover, O. 2011. The Spanish Survey of Household Finances (EFF): Description and methods of the 2008 wave. Banco de España Occasional Paper 1103.

Bricker, J., A. B. Kennickell, K. B. Moore and J. Sabelhaus. 2012. Changes in U.S. family finances from 2007 to 2010: Evidence from the Survey of Consumer Finances. In: Federal Reserve Bulletin 98(2). 1–80.

Cameron, A. and P. Trivedi. 2005. Microeconometrics: Methods and applications. Cambridge University Press.

Cochran, W. G. 1977. Sampling techniques. 3rd edition. Wiley.

Cowles, M. K. and B. P. Carlin. 1996. Markov chain Monte Carlo convergence diagnostics: A comparative review. In: Journal of the American Statistical Association 91(434). June. 883–904.

Drescher, K., P. Fessler and P. Lindner. 2020. Helicopter money in Europe: New evidence on the marginal propensity to consume across European households. Economics Letters, Volume 195. October. https://www.sciencedirect.com/science/article/abs/pii/S0165176520302603 (accessed on October 2, 2025).

ECB. 2013a. The Eurosystem Household Finance and Consumption Survey: Methodological report for the first wave. ECB Statistics Paper Series 1. April.

ECB. 2013b. The Eurosystem Household Finance and Consumption Survey: Results from the first wave. ECB Statistics Paper Series 2. April.

ECB. 2016. The Eurosystem Household Finance and Consumption Survey: Methodological report for the second wave. ECB Statistics Paper 17. December.

Fessler, P. and M. Schürz. 2013. Cross-country comparability of the Eurosystem Household Finance and Consumption Survey. In: Monetary Policy & the Economy Q2/13. OeNB. 29–50.

Fessler, P. and M. Schürz. 2015. Private wealth across European countries: The role of income, inheritance and the welfare state. ECB Working Paper Series 1847.

Fessler, P. and M. Schürz. 2017. Zur Verteilung der Sparquoten in Österreich. In: Monetary Policy & the Economy Q3/17. OeNB. 13–33.

Fessler, P. and M. Schürz. 2019. Vermögen der privaten Haushalte in Österreich - Gemeinsamkeiten und Unterschiede. In: Soziale Mobilität und Vermögensverteilung (BMASK). 71–92.

Fessler, P. and M. Schürz. 2022. Structuring the Analysis of Wealth Inequality Using the Functions of Wealth: A Class-Based Approach. In: Measuring Distribution and Mobility of Income and Wealth. National Bureau of Economic Research. October. https://www.nber.org/books-and-chapters/measuring-distribution-and-mobility-income-and-wealth/structuring-analysis-wealth-inequality-using-functions-wealth-class-based-approach (accessed October 2, 2025).

Fessler, P., K. Jäger-Gyovai and T. Messner. 2015. What can we learn from Eurosystem Household Finance and Consumption Survey data? An application to household debt in Slovakia. In: Focus on European Economic Integration Q2/15. OeNB. 76–87.

Fessler, P., P. Lindner and E. Segalla. 2014. Net wealth across the euro area – why household structure matters and how to control for it. ECB Working Paper 1663.

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. OeNB. 35–96.

Fessler, P. , P. Lindner and M. Schürz, 2019. Eurosystem Household Finance and Consumption Survey 2017 for Austria. In: Monetary Policy & the Economy Q4/18. OeNB. 36–66.

Fessler, P., E. List and T. Messner. 2017. How financially vulnerable are CESEE households? An Austrian perspective on its neighbors. In: Focus on European Economic Integration Q2/17. OeNB. 58–79.

Fessler, P., P. Mooslechner and M. Schürz. 2012. Eurosystem Household Finance and Consumption Survey 2010: First results for Austria. In: Monetary Policy & the Economy Q3/12. OeNB. 23–62.

Frumento, P., F. Mealli, B. Pacini and D. B. Rubin. 2012. Evaluating the effect of training on wages in the presence of noncompliance, nonemployment, and missing outcome data. In: Journal of the American Statistical Association 107(498). 450–66.

Haslinger, A. and J. Kytir. 2006. Stichprobendesign, Stichprobenziehung und Hochrechnung des Mikrozensus ab 2004. In: Statistische Nachrichten 6/2006. 510–19.

Haziza, D. and J.-F. Beaumont. 2007. On the construction of imputation classes in surveys. In: International Statistical Review 75. 25–43.

Heeringa, S. G., B. T. West and P. A. Berglund. 2017. Applied survey data analysis. 2[nd] edition. Chapman Hall/ CRC Press.

Iannacchione, V. G., J. G. Milne and R. E. Folsom. 1991. Response probability weight adjustments using logistic regression. In: Proceedings of the American Statistical Associations, Section on Survey Methods. 637–42.

Kennickell, A. B. 1998. Multiple imputation in the Survey of Consumer Finances. In: Proceedings of the Section on Business and Economics Statistics. 1998 Annual Meetings of the American Statistical Association. 63–74.

Kennickell, A. B. 2005. The good shepherd: Sample design and control for wealth measurement in the Survey of Consumer Finances. Federal Reserve Board. January.

Kennickell, A. B. 2011. Look again: editing and imputation of SCF panel data. Prepared for the Joint Statistical Meeting in Miami, Florida. August 3.

Kennickell, A. B., P. Lindner and M. Schürz. 2021. A new instrument to measure wealth inequality: distributional wealth accounts. In: Monetary Policy & the Economy Q4/21. OeNB. https://www.oenb.at/dam/jcr:37664c81-2d0d-409e-8d33-a19fc2b25854/05_mop_q4_21_A-new-instrument-to-measure-wealth-inequality.pdf (accessed on October 2, 2025).

Kennickell, A. B. and D. McManus. 1993. Sampling for household financial characteristics using frame information on past income. In: Proceedings of Survey Research Methods. Section of the American Statistical Association. 88–97.

Kish, L. 1995. Survey sampling. Wiley.

Kolenikov, S. 2010. Resampling variance estimation for complex survey data. In: The Stata Journal 10(2). 165–99.

Kovar, J. G., J. N. K. Rao and C. F. J. Wu. 1988. Bootstrap and other methods to measure errors in survey estimates. In: The Canadian Journal of Statistics 16. 25–45.

Levy, P. and S. Lemeshow. 2008. Sampling of populations: Methods and applications. 4th edition. Wiley.

Lindner, P. 2021. Finanzvermögen der privaten Haushalte aus Perspektive der Mikrodaten. In: Statistiken Sonderheft: Einkommen, Konsum und Vermögen der Haushalte – Sektorale Volkswirtschaftliche Gesamtrechnungen in den letzten 20 Jahren. OeNB. https://www.oenb.at/dam/jcr:5ff19df4-e600-43ec-9782-915b7dce16ba/05_SH_Sektorale-VGR_2021_Lindner.pdf (accessed on October 2, 2025).

Lindner, P. and M. Propst. 2020. Interviewdauer des HFCS in Österreich. In: Statistiken - Daten und Analysen Q2/20. OeNB. https://www.oenb.at/dam/jcr:4d0725ce-c21d-45e3-9405-886f1d55ca2d/08_statistiken_Q2_20_Interviewdauer-des-HFCS.pdf (accessed on October 2, 2025).

Lindner, P. and M. Schürz. 2015. Varianten der Messung von Haushaltsvermögen im HFCS in Österreich. In: Statistiken – Daten und Analysen Q2/15. OeNB. 52–70.

Lindner, P. and M. Schürz. 2017. Kommentare von Respondenten des Household Finance and Consumption Survey zur Befragung. Statistiken - Daten und Analysen Q4/17. OeNB. 50–63.

Lindner, P. and M. Schürz. 2019. The joint distribution of wealth, income and consumption in Austria: a cautionary note on heterogeneity. In: Monetary Policy & the Economy Q4/19. OeNB. https://www.oenb.at/dam/jcr:ed43b82c-cbaa-45db-bce2-07d4e1d1faf9/joint_distribution_mop_%-20q4_19_screen-5.pdf (accessed on October 2, 2025).

Lindner, P. and M. Schürz. 2021. Matching survey data on wealth to register data on pension entitlements: what challenges need to be addressed? In: Statistiken – Daten und Analysen Q3/21. OeNB. https://www.oenb.at/dam/jcr:d4d8b02f-8a0e-4dac-82cd-89e96709952e/08_PB_statistiken_Q3_21_Matching%20survey%20data.pdf (accessed on October 2, 2025).

Lindner, P. and V. Redak. 2017. The resilience of households in bank bail-ins. In: Financial Stability Report 33. OeNB. 88–101.

Lindner, P., M. Schürz and J. C. Zhan. 2014. Methodische Verbesserungen im HFCS Austria. In: Statistiken – Daten und Analysen Q4/14. OeNB. 71–83.

Lindner, P., T. Y. Mathä., G. Pulina and M. Ziegelmeyer. 2022. Borrowing constraints, own labour and homeownership. Applied Economics, 55(42), 4931–4945. https://doi.org/10.1080/00036846.2022.2133893 (accessed on October 2, 2025).

Little, R. J. A. and D. B. Rubin. 2019. Statistical analysis with missing data. 3rd edition. Wiley Series in Probability and Statistics.

Little, R. J. A. and S. Vartivarian. 2003. On weighting the rates in non-response weights. In: Statistics in Medicine 22(9). 1589–99.

OeNB. 1979. Mitteilungen des Direktoriums der Oesterreichischen Nationalbank.

OeNB. 1998. Statistisches Monatsheft. December.

Pattengale, N. D., M. Alipour, O. R. P. Bininda-Edmonds, B. M. E. Moret and A. Stamatakis. 2010. How many bootstrap replicates are necessary? In: Journal of Computational Biology 17(3). 337–54.

Rao, J. N. K. and C. F. J. Wu. 1988. Resampling inference with complex survey data. In: Journal of the American Statistical Association 83. 231–41.

Rao, J. N. K., C. F. J. Wu and K. Yue. 1992. Some recent work on resampling methods for complex surveys. In: Survey Methodology 18. 209–17.

Royston, P. 2004. Multiple imputation of missing values. In: Stata Journal 4(3). 227–41.

Schafer, J. L. and M. K. Olsen. 1998. Multiple imputation for multivariate missing data problems: A data analyst's perspective. In: Multivariate Behavioral Research 33. 545–71.

Statistics Austria. 2013. Bevölkerungsstand 1.1.2013.

Statistics Austria. 2018. Standard-Dokumentation Metainformationen (Definitionen, Erläuterungen, Methoden, Qualität) zu EU-SILC 2017.

Valliant, R., J. A. Dever and F. Kreuter. 2013. Practical tools for designing and weighting survey samples. Springer.

Van Buuren, S. and C. G. M. Oudshoorn. 1999. Flexible multivariate imputation by MICE. TNO-rapport PG 99.054. TNO Prevention and Health. Leiden.

Van Buuren, S., H. C. Boshuizen and D. L. Knook. 1999. Multiple imputation of missing blood pressure covariates in survival analysis. In: Statistics in Medicine 18(6). 681–94.

Van Buuren, S., J. P. Brand, C. G. Groothuis-Oudshoorn and D. B. Rubin. 2006. Fully conditional specification in multivariate imputation. In: Journal of Statistical Computation and Simulation 76(12). 1049–64.

Vehovar, V. 1999. Field substitution and unit nonresponse. In: Journal of Official Statistics 15. 335–50.

Wagner, K. 2014. Intergenerational transmission: How strong is the effect of parental homeownership? In: Monetary Policy & the Economy Q2/14. OeNB. 49–64.

Williams, R. L. 2014. Survey Sampling and Weighting, In: Encyclopedia of Health Economics. Culyer, A. J. (ed.). Elsevier. 371–74.