

MONETARY POLICY & THE ECONOMY

Quarterly Review of Economic Policy

Eurosystem Household Finance and
Consumption Survey 2014

Methodological notes for Austria
(second wave)

Contents

Key facts of the HFCS in Austria	4
1 Introduction	7
2 Questionnaire	9
2.1 Introduction	9
2.2 Objectives of the survey	9
2.3 Unit of collection	10
2.4 Data collection period and reference period	11
2.5 Interview structure and content	11
2.6 Special features	17
2.7 Interviewer documents	21
2.8 Participating countries	21
2.9 Online appendix	22
3 Interviewers	23
3.1 The interviewers' role in the survey process	23
3.2 General information	23
3.3 Interviewer training	23
3.4 Contact strategies and specifications	26
3.5 Documents and other supporting material	26
3.6 Monitoring	28
3.7 Problems relating to interviewers	29
3.8 Survey of interviewers	29
3.9 Online appendix	30
4 Consistency checks and editing	31
4.1 Introduction	31
4.2 Number and type of edits	31
4.3 Consistency checks during interviews	32
4.4 Postinterview consistency checks	33
4.5 Flags	35
4.6 Ex post editing	38
4.7 Formatting and editing after multiple imputations	48
4.8 Concluding remarks and online appendix	50
5 Multiple imputations	51
5.1 Introduction	51
5.2 Item nonresponse	51
5.3 HFCS imputation procedure	54
5.4 Creating the imputations	55
5.5 Selected results	62
5.6 Concluding remarks	64

6	Sampling	65
6.1	Introduction	65
6.2	Target population and sampling frame	65
6.3	Background – the (external) datasets used	67
6.4	Stratification and sample size	69
6.5	The two stages of the random draw	72
6.6	Concluding remarks	73
7	Construction of survey weights	75
7.1	Introduction	75
7.2	Construction of survey weights	76
7.3	Selected results	84
7.4	Concluding remarks	85
8	Construction of replicate weights for variance estimation	86
8.1	Introduction	86
8.2	Construction of replicate weights	87
8.3	Concluding remarks	89
9	User guide	91
9.1	Introduction	91
9.2	Merging the data files	91
9.3	Multiple imputations	93
9.4	Survey variables	95
9.5	Standard estimation procedures	96
9.6	Additional estimation procedures	97
9.7	Online appendix	97
10	Changes from the first to the second wave of the HFCS	98
10.1	Introduction	98
10.2	Questionnaire	98
10.3	Interviewers	101
10.4	Consistency checks and editing	102
10.5	Multiple imputations	102
10.6	Sampling	103
10.7	Construction of survey weights	104
10.8	Construction of replicate weights for variance estimation	106
10.9	Concluding remarks	106
	References	107

Opinions expressed by the authors do not necessarily reflect the official viewpoint of the Oesterreichische Nationalbank or of the Eurosystem.

Key facts of the HFCS in Austria

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, integrating 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 June 2014 and February 2015). Income-related data refer to the 2013 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 (“Zählsprenkel”)
- Secondary sampling unit (SSU): postal addresses

The gross sample comprised a total of 619 PSUs and 6,308 SSUs in 185 strata.

Survey company

Institut für empirische Sozialforschung GmbH – IFES

Fieldwork

General information

Fieldwork period:	June 2014 to February 2015
Number of interviewers:	72
Method of data collection:	computer-assisted personal interview (CAPI)

Interviewer training

Number of interviewer training sessions:	6
Duration of HFCS interviewer training:	1 day

Pilot survey

Number of pilot interviews:	55
-----------------------------	----

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 silver coin with a face value of EUR 5 (worth some EUR 15 at the time of fieldwork).

In addition, all participating households had the opportunity to enter a draw for travel vouchers (one worth EUR 1,000, five vouchers worth EUR 200 each).

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 every one in six interviews.

The (anonymized) data from the completed household interviews were forwarded to the OeNB in 15 batches during the field phase, to enable prompt assessment of each interview and interviewer. Four interviewers who were found to have delivered seriously flawed interviews were withdrawn from the survey.

Follow-up queries by telephone

Outliers and inconsistent responses were looked into again after the interviews. These follow-up investigations resulted in the confirmation or correction of the respective values (for 370 households).

Editing and consistency checks

Number and type of edits

Number of observations:	around 1.3 million
Thereof edited observations:	around 65,000 (more than two-thirds on the basis of verbatim records)
Percentage of edited observations:	4.8%

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.

Documentation

Flag variables are used to document all edits and imputations.

Imputations

Method:	multiple imputation by 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:	18
Mean of the variables with missing values per household:	29.9

Sample size and response rate

Number of households in the sample (gross sample):	6,308
Number of successfully interviewed households (net sample):	2,997
Number of households that could not be reached (despite five contact attempts):	136
Number of households that refused to participate:	2,657
Number of households that did not participate for other reasons:	204
Number of addresses whose eligibility was unknown:	30
Number of ineligible addresses:	284
Incomplete interviews and interviews discarded after fieldwork:	42
Response rate:	around 50%

Weighting

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

- Method of nonresponse adjustment: model-based adjustment combined with weighting-class adjustment, based on optimum number of classes
- Method of poststratification adjustment: cell adjustment

Smallest final weight:	287
Median final weight:	1,207
Mean final weight:	1,289
Highest final weight:	4,360
Sum of final weights (target population):	3,862,526
Unequal weighting effect:	1.167
Trimming and normalization of weights	none

Variance estimation

Method:	rescaling bootstrap procedure
Number of replicates:	1,000
Number of pseudo-strata:	135
Replicate weights computed with same adjustments to design weights as those to the final weights	
Finite population corrections were applied to all resampling 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. In Austria, the euro area HFCS has so far been carried out twice: in 2010/11 (first wave) and in 2014/15 (second wave). At the euro area level, the geographical scope of the HFCS was broadened during the second wave to include three countries which did not participate in the first wave, namely Ireland, Estonia and Latvia. HFCS data are comparable across all euro area countries thanks to ex ante harmonization of the survey and of the survey methods. In Austria, the HFCS was conducted by the Oesterreichische Nationalbank (OeNB) in cooperation with the survey company IFES (Institut für empirische Sozialforschung GmbH). The ECB is scheduled to make data for all euro area countries from the second HFCS wave available for research purposes from fall 2016.

This publication, entitled *Eurosystem Household Finance and Consumption Survey 2014: Methodological notes for Austria (second wave)*, provides an in-depth view of the data collection process and the methods applied. Based on the methodological documentation from the first HFCS wave in Austria (Albacete et al., 2012), it aims at making the process of data collection as transparent as possible and serves as the basis for correct evaluation of HFCS data. Between the two waves, specific methodological aspects were discussed in a number of publications. For instance, the information gathered from the survey of interviewers has been examined in depth and cross-checked with the HFCS data (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. computing survey weights and producing correct variance estimates from HFCS data) are arranged in a way to ensure comprehensibility. To avoid redundancies, only essential details were repeated. The following eight chapters provide a detailed explanation of each step in the survey process, with another chapter summarizing the changes compared to the first wave.

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 qualification of the HFCS interviewers, as their professional demeanor and expertise contribute 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 queries 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 have been flagged, like all edits. Users of the HFCS data may apply our imputations or deal with item nonresponse in a different way. The complex survey sample design used for the first wave was enhanced to ensure a sufficiently representative sample of Austrian households that fits the purpose of the Eurosystem and the OeNB. Chapter 6 on *Sampling* provides a detailed description of this survey sample design.

The final household weights were calculated in several steps on this basis. Chapter 7 outlines the procedure for *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 correct variance estimates, which is presented in chapter 8 on *Construction of replicate weights for variance estimation*.

The *User guide* in chapter 9 provides basic advice on the correct use of HFCS data in Stata®.

Experienced HFCS users will find a summary of all material enhancements and adaptations implemented in the second wave in chapter 10 (*Changes from the first to the second wave of the HFCS*). This comparative section refers back to and expands on the detailed documentation provided in the previous sections.

All the essential documentation and background material used in the HFCS are published in the *Online appendix* at www.hfcs.at/en. This website also provides information about the publication of HFCS data from all participating countries by the ECB (expected in fall 2016) and any other HFCS news.

2 Questionnaire

2.1 Introduction

In December 2006, the Governing Council of the 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. Despite cross-country differences in the technical conditions for conducting the Eurosystem HFCS, for example with respect to sampling design (chapter 6) and multiple imputation (chapter 5), ex ante harmonization has been achieved with regard to the information to be collected, and broadly achieved with regard to methodology. The plan is to conduct the HFCS every three years; in Austria, it has so far been conducted twice (HFCS 2010 and HFCS 2014).

The questionnaire used for the second HFCS wave was modeled on that of the first wave. This chapter presents the Austrian questionnaire, which was designed on the basis of the euro area blueprint questionnaire (drafted in English).¹ The Austrian survey covers both the internationally agreed variables and country-specific features (e.g. foreign currency loans 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) and define the unit of collection (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 other 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 euro area households' assets and liabilities, as these data allow for the analysis of households' investment and consumption decisions. These survey data can be used e.g. 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.

The HFCS, which is the most comprehensive household-level survey conducted on this subject in the euro area, has now been carried out a second time.²

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. Also, as the economic developments of the past decade have clearly shown, 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 analysis of the specific burden on different income, occupational and age groups. Hence,

¹ For further details about the HFCN and the HFCS (including the euro area blueprint questionnaire), see www.ecb.europa.eu/pub/economic-research/research-networks/html/researcher_hfcn.en.html (accessed on December 9, 2016).

² Some euro area countries (such as Ireland and Estonia), which did not participate in the first wave (HFCS 2010), have conducted the HFCS only once so far.

decision-making on monetary policy and financial stability issues is also informed by HFCS-based analyses. Publications attesting to the variety of insights gained from analyses relying on the microdata from the first HFCS wave are, amongst others, Albacete and Lindner (2013 and 2015), Albacete et al. (2014), Fessler and Schürz (2013), Fessler et al. (2015) and Wagner (2014).

2.3 Unit of collection

To begin with, designing a survey involves defining the target population (see also chapter 6) and the unit of data collection. In the case of the HFCS, households represent the main unit of collection, but some data are also collected at the individual level.

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. 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 and remain financially dependent on the household.

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 single address can 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

2.3.2 Financially knowledgeable person

All household questions 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

³ See chapter 6 for a more detailed definition of a household.

person (FKP), answered all questions relating to the household as a whole (green sections in chart 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, he or she was also the reference person of the given household. While this continued to be the case with almost all households, it was not a prerequisite in the second wave of the HFCS in Austria. The FKP could also be a family member (e.g. a son or daughter) who was in charge of the household's finances but was no longer a member of the household. Theoretically, a household's tax consultant or financial advisor could also participate as the FKP in the HFCS on behalf of the household, but there was no such case in Austria. In cases where the FKP was not a member of the household, a reference person was selected from the members of 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 June 2014 to February 2015. In contrast, the questions about income (except the one about average monthly net household income)⁴ referred to the calendar year 2013 (reference period), as this was the last full calendar year before the survey was started.

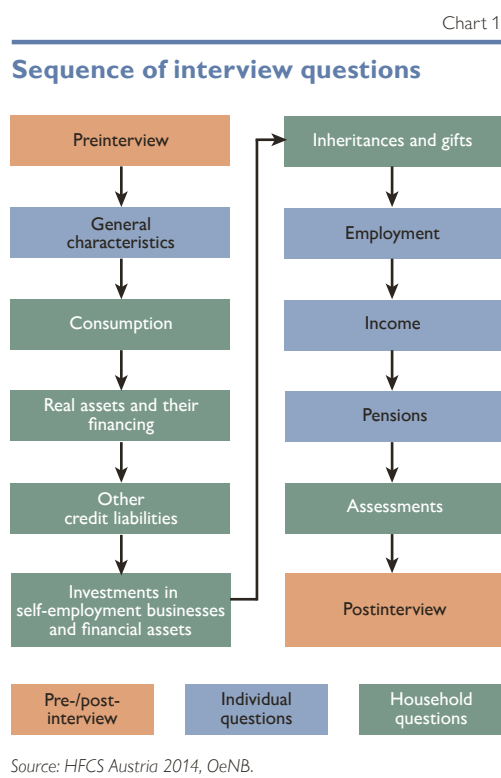
2.5 Interview structure and content

2.5.1 Questionnaire structure

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

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,



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

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

in line with the HFCS definition, listing the individual household members and identifying the main respondent among them, i.e. the FKP or, if the FKP was not a household member, a reference person. This was also the stage when key identifying data were collected for all household members, namely gender, age and relation to the main respondent. Finally, all 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 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, and, if applicable, migration background. 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) served 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, total consumption spending and transfers to people outside the household. The information on household income was collected to verify whether this income was high enough to finance the household's expenditure; if not, further information was collected on how households financed any expenses in excess of their income.

Real assets and their financing

This section contained questions about the household's housing situation and most of the other real assets (excluding investments in self-employment businesses, which are addressed below) 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.

Homeowners were asked to indicate when and how they had acquired their main residence as well as how much the property was worth at the time of the interview and at the time they first acquired ownership. In addition, homeowners were asked whether they had taken out a mortgage on their main residence to finance the property. For outstanding mortgages using the residence as collateral, the following data were collected separately for a maximum of three loans: initial amount, 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.

Households living in a rented home were asked about their rental cost (including and excluding utilities). 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 together with unsecured loans (see “Other credit liabilities”).

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, as highlighted in the questionnaire. Specifically, loop questions⁶ were used to identify the fundamentals (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, and a summary question was used to establish the total value for any further properties. Any outstanding mortgages taken out on these properties were established using loop questions for a maximum of three loans, as with the main residence. Again, only a summary of the information was collected for additional mortgages taken out on additional properties.

Finally, households were asked to indicate the value of any cars or other vehicles as well as of any valuables (e.g. jewelry, 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 credit liabilities

This category covered all other credit liabilities: leasing contracts, outstanding balances on current accounts, outstanding balances on credit cards and uncollateralized 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 monthly amount of lease payments instead of the total amount outstanding). For outstanding balances on credit cards, the interviewer asked only about the amount in excess of the usual payments, not about the monthly bill that is repaid at the end of a particular month.

The questions about unsecured loans were asked in two parts using a loop with up to three iterations. First, households were asked whether they had also borrowed money from relatives or friends, in addition to the loans already recorded, and they were asked to indicate the purpose and outstanding amount of any such loans. All other (unsecured) loans were recorded in a second part. If a household held more than three such loans, interviewers only recorded the total outstanding principal and, for other credit liabilities, the sum of repayment rates. In the loop for other loans, respondents were asked about the same indicators as for the mortgage(s) secured against property.

In addition, this section of the questionnaire addressed households’ risk attitude and the outcome of any loan application they might have made recently.

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

Private businesses and financial assets

In this section, the HFCS survey documented all wealth components of the household balance sheet beyond the real assets covered before. Interviewers started out by asking whether a household, partly or entirely, owned any self-employment businesses, and, if so, recorded separate information for up to three such businesses (industry, legal form, number of employees and current value). Information about any additional private businesses was summarized to provide a total value.

The next questions focused on assets held in sight accounts, savings accounts, savings plans with building and loan associations, life insurance funds, mutual funds, bonds, listed stock, silent partnerships, 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 all life insurance contracts, interviewers asked about the date of conclusion, the type (benefits to be provided at the death of the policy holder or at a given date, or a combination thereof), duration of contract and the frequency and amount of payments into the life insurance plan, to enable a projection of the amounts held in such funds. In addition, interviewers asked about financial debt owed to the household and other financial assets as well as the household's estimated total wealth. This estimate was used to assess the plausibility⁷ 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. Only the total value was recorded for all inheritances or gifts in excess of five. The dataset also includes information about inheritances or gifts households expect to receive.

Up to this point, all questions (excluding those about the sociodemographic makeup of household) focused on the household as a whole. In contrast, the questions in the three subsequent sections of the HFCS questionnaire related to individual household members aged 16 and over, and not to the household as a whole.

Employment

Household members aged 16 and over were asked to provide information on their employment status. The first set of questions focused on people active in the labor market. Pensioners, homemakers, schoolchildren, 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

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

⁸ The international core dataset only contains up to three inheritances/gifts. Information about more than three inheritances/gifts is not provided. More than three inheritances were reported by 0.3% of households.

code⁹), the number of working hours (and seasonal fluctuations), the employer's main economic activity (ÖNACE code¹⁰), 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

The 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 2013, as this was the last full calendar year before the start of fieldwork in 2014.

The types of income covered were employee and self-employment income, income from the state pension system and from private and occupational 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.)

In addition to asking about income at the individual level, the interviewer also asked in the interview with FKPs whether a household received income from regular (public or private) social transfers, from the rental or leasing of 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.

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 occupational or private pension plans.

Assessments¹²

In this section, FKPs were invited to provide their assessment of social issues in Austria and of their households' position in the national wealth distribution.

⁹ ISCO: *International Standard Classification of Occupations*, see www.ilo.org/public/english/bureau/stat/isco/isco08/index.htm (accessed on December 9, 2016).

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 www.statistik.at/web_en/classifications/implementation_of_the_onace2008/index.html (accessed on December 9, 2016).

¹¹ *The questions from the second part of this section in the questionnaire are not part of the international core dataset.*

¹² *These are noncore variables specific to Austria that is not included in the international HFCS dataset.*

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 and any other issues. 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 IBM's Quancept CAPI software on the basis of a Word template, which is available in the online appendix. In addition to filtering, the questionnaire also contains internal consistency checks (see chapter 4) that allow corrections of possible errors in data entry while the interview is ongoing. The use of the CAPI (computer-assisted personal interviewing) technique combines the advantages of a personal interview with those of real-time digital recording and data cross-checking. In addition, it allows for the implementation of complex filtering techniques, producing tailor-made 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 55 households.

2.5.2.3 CAPI problems

Given the experience gained with the first wave of the HFCS in Austria and in-depth checks of the questionnaire and of the questionnaire programming before the start of fieldwork, it was possible to eliminate errors to a great extent. Below is a list of reprogramming issues that had to be dealt with during the field phase:

- Account balances of public pension plans (PF0510 and APF0520):¹³ In the initial field phase there was a problem with filtering based on variable PF0300 (number of contribution years) which caused "Don't know" respondents to be wrongly routed to question PF0600 (occupational pension), skipping the variables APF0520 and PF0510. Only 13 people were affected by this problem. The missing information was queried by phone, and the dataset was corrected.
- Confirmation of numeric variables (also section 2.6.2 on euro loops): Again in the initial field phase two households had indicated negative net wealth. In the case of these households, the minus sign was also indicated by mistake in the subsequent cross-checking of the amounts, even though the amounts recorded were positive. The information provided by those households was queried, and the questionnaire programming was corrected.
- Additional borrowing (HB150\$x): This variable was incorrectly filtered for a couple of weeks during fieldwork, causing this question to be skipped. This glitch was corrected after a total of 39 households had not answered this question by mistake. The missing data were flagged for imputation (see also chapter 5), to put no unnecessary burden on the respondents.

¹³ This is a noncore variable specific to Austria that is not included in the international HFCS dataset.

- Checking the consistency of questions on income from dependent employment and the public pension system (PG0100 and PG0300): During the field phase, an additional consistency check was added for these two questions (see also section 4.3). The consistency check served to alert respondents to an unusual mix of answers if they had first indicated to have been dependently employed or drawing retirement benefits for several years, but subsequently answered “No” to the yes/no question on the relevant type of income (income from dependent employment or the public pension system).

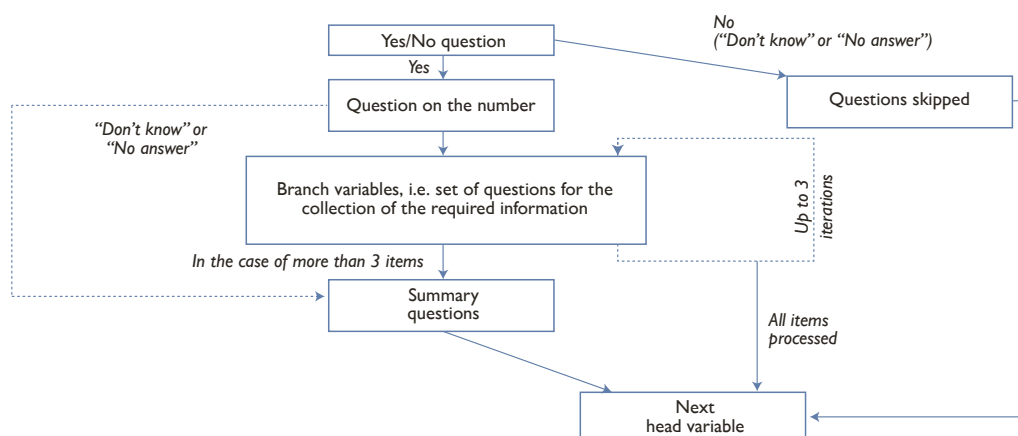
2.6 Special features

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 each of possibly multiple items applicable to a household (e.g. loans). Chart 2 shows the schematic cycle of these loops.

Chart 2

Structure of a loop



Source: HFCS Austria 2014, OeNB.

Loops were used for the following items:

- mortgage loans using the household’s main residence as collateral
- real estate holdings other than the household’s main residence
- mortgage loans that used these other properties as collateral
- unsecured loans from family and friends
- other unsecured loans
- investments in self-employment businesses (other than holdings of listed shares)
- life insurance contracts (no maximum number of iterations)
- ownership transfers through inheritances or gifts

First, respondents indicated whether a specific item applied to their household, and if so, the number of such items held by the household. Then the details were recorded in a loop for every one of these items. 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. As a rule, this procedure was repeated up to three times. If a household had received

multiple inheritances and gifts, the procedure was repeated up to five times. When there were more than three items (five in the case of gifts and inheritances), interviewers recorded a summary of the information.¹⁴ The sequence of items was predefined in the questionnaire. In the case of life insurance contracts, the questionnaire did not predefine a maximum number of iterations in order to record all the contracts.

To make the survey as user-friendly as possible, interviewers could exit loops 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.

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 the so-called euro loops. See the online appendix on “Euro loops” for a schematic overview.

In the first step, respondents were asked to provide specific amounts (“How much...?”). If the amount was stated, the respondent could freely select any 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 as before when seeking to get specific amounts, asking respondents to first indicate the currency and then confirm the range and the currency.

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 the questions in the first wave 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 repayment of loans. 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.¹⁵ The predefined ranges referred to amounts in euro

¹⁴ Only very few households indicated that they held more than three of the relevant items.

¹⁵ See the online appendix at <https://hfcs.at/en/publikationen/dokumentation.html> for a detailed overview of which ranges (list A, B or C) were used for which questions.

Table 1

List of predefined ranges for euro questions

List A	List B	List C
<i>EUR</i>		
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 million	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 2014, OeNB.

only. Interviewers moved on to the confirmation question as soon as the respondent had chosen a range. The lists of predefined ranges were used as 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 on ranges was especially important for multiple imputation (see chapter 5).

2.6.3 Recording households owning a farm

Recording the real assets of farmers is known – from past experience of the first wave – to be a particular challenge for respondents, especially when it came to breaking down the assets into the household main residence and business assets. In the second wave, the business assets of farmers were again recorded in the loop on investment in self-employment businesses. However, in order to elicit more precise answers, the second-wave questionnaire contained a number of additional questions for farmers as well as some additional guidance. The procedure can be summarized as follows:

- Before the actual interview started, respondents were classified by interviewers as farmers or non-farmers. The classification was straightforward in almost all cases. Even when a respondent was incorrectly identified, all essential information was still gathered.
- 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, farmers 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 Additional information

2.6.4.1 Contact attempts

Every household in the sample population had to be contacted unsuccessfully on at least five separate occasions before it could be classified as a unit nonresponse (see also the contact rules in section 3.4).¹⁶

The contact attempts were recorded in the dataset,¹⁷ 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 workload (interview, refusal, etc.) on a particular household.

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 venue 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

¹⁶ Interviewers had to make at least two contact attempts in person over a period of at least three weeks.

¹⁷ These variables were not included in the user database due to anonymization requirements.

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.¹⁸ The second section (excluding interviewer comments) was recorded in variables HR0100 to HR1600 in the household data file.

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):¹⁹

- the showcards which were used during interviews to provide respondents with a list of possible 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 “Eurosysteem Finance and Consumption Survey 2010: First Results for Austria” (Fessler et al., 2012) to illustrate how the data obtained in the first wave of the HFCS in Austria were used for analysis.

2.8 Participating countries

The second wave of the Eurosystem HFCS was conducted in the following euro area countries: Belgium, Germany,²⁰ Estonia, Ireland,²¹ Greece, Spain,²² France, Italy,²³ Cyprus,²⁴ Luxembourg, Latvia, Malta,²⁵ Austria, the Netherlands, Portugal, Slovenia,²⁶ Slovakia and Finland.

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

¹⁸ 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).

¹⁹ For a detailed description of the documentation, see chapter 3.

²⁰ Information on the survey in Germany is available at www.bundesbank.de/Navigation/EN/Bundesbank/Research/Panel_on_household_finances/panel_on_household_finances.html (accessed on December 9, 2016).

²¹ Information on the survey in Ireland is available at www.cso.ie/en/surveysandmethodology/socialconditions/hfcsurvey/ (accessed on December 9, 2016).

²² Information on the survey in Spain is available at www.bde.es/bde/en/areas/estadis/Otras_estadistic/Encuesta_Financi/ (accessed on December 9, 2016).

²³ Information on the survey in Italy is available at www.bancaditalia.it/statistiche/tematiche/indagini-famiglie-imprese/bilanci-famiglie/index.html?com.dotmarketing.htmlpage.language=1 (accessed on December 9, 2016).

²⁴ Information on the survey in Cyprus is available at www.centralbank.gov.cy/nqcontent.cfm?a_id=11840&lang=en (accessed on December 9, 2016).

²⁵ Information on the survey in Malta is available at <https://www.centralbankmalta.org/en/household-finance-and-consumption-survey> (accessed on December 9, 2016).

²⁶ Information on the survey in Slovenia is available at www.bsi.si/en/financial-data.asp?MapaId=1584 (accessed on December 9, 2016).

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 in addition to the datasets that the ECB is expected to publish in fall 2016.

2.9 Online appendix

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

- the questionnaire
- the euro loops
- the paradata questions
- the variable list
- the showcards
- the glossary

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

A total of 72 interviewers were involved in the second wave of the HFCS. 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 to have specific experience, either by having been involved in the first wave 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) to qualify for the job. In fact, about half of the interviewers in the second wave of the HFCS in Austria also conducted interviews in the first wave. Payment for successfully completed interviews was calculated on the basis of the surveyed household's 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 unsuccessful interviews, the interviewers were required to have made at least two personal contact attempts and five contact attempts altogether.

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 survey company organized a total of six one-day training sessions with an expert from the HFCS team (OeNB staff member), which took place in Vienna (twice), Graz, Linz, Innsbruck and Zams before the start of the fieldwork. Only trained interviewers were used in the HFCS. Contrary to the first HFCS wave, the survey company did not hold additional training sessions for newly hired interviewers.

A training session consisted of an all-day (10:00 a.m. to 6:30 p.m.) interactive workshop, in which the interviewers were encouraged to ask questions as they arose. Essentially, the workshop focused on four main areas, as described below.

3.3.1 Training unit 1

Introduction

First, a member of the OeNB's 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 the 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, which was split into three parts, 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. In this practical example, the interviewer asked the questions and the lecturer provided the answers. This approach made it possible for participants to acquaint themselves with the essential elements of the questionnaire both in theory and practice.

Questionnaire – theory and practice 1

The first part of this training unit covered the preinterview questionnaire including the recording of a household in the 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 and record farmers were also given ample time. This part closed with a practical conduction of an interview from the beginning of the questionnaire to the end of the section on households' real estate wealth.

Questionnaire – theory and practice 2

Part two of this training unit covered the treatment of other liabilities, private businesses and financial assets, as well as the section on inheritances and gifts. 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. At the end of this part of the training unit, the practical interview started in part one was continued to the end of the section on inheritances and gifts.

Questionnaire – theory and practice 3

The third and last part of this training unit was dedicated to the individual data to be collected in the interview, namely information on household members' employment status, income and retirement provision. Moreover, this training unit covered 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 the comments provided by the respondents before concluding this session with the interview.

3.3.3 Training unit 3

Interviewers' tasks, contact provisions and paradata

One of the interviewers' central tasks was to convince the selected households to take part in the HFCS. 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 of 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 must document their attempts to contact the selected household and record 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 an as neutral way as possible (using the glossary, if necessary). Comments received from the first wave of the HFCS in Austria helped to identify and discuss interview situations that are commonly encountered.

3.3.4 Training unit 4

Documents and other material

In training unit 4, lecturers and interviewers once again went through and explained all the documentation and material made available (see section 3.5) to the interviewers, which had been used in training units 1 to 3. 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 disbursement of household addresses that they had to contact. Also, they received information about the incentives for households that completed an interview and about interviewers' remuneration.

3.4 Contact strategies and provisions

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 (five weeks in August). 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 to households

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 silver coin with a face value of EUR 5 (worth some EUR 15 at the time of fieldwork). The interviewers handed over the silver coins to the respondents directly upon completion of the questionnaire. In addition, each household had the opportunity to take part in a lottery drawing for one travel voucher worth EUR 1,000 and five further travel vouchers worth EUR 200 each.

3.5.3 Scientific study

The interviewers were instructed to have with them a copy of the study “Eurosystem Household Finance and Consumption Survey 2010: First Results for Austria” by Fessler et al. (2012) (see online appendix) during each contact attempt. This study is based on data taken from the first HFCS wave in Austria and gives 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. Feedback by the interviewers after the first wave showed that reluctant respondents in particular were more likely to participate in the survey after having received this information.

3.5.4 Catalogue of 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
- types of income
- loan repayment
- lending institution
- economic sectors
- types of life insurance contracts
- types of mutual funds
- banks
- investment behavior
- type of inheritance/gift
- employment status I and II
- smoking habits
- family background – consumption
- topics

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 the particular places where the use of a showcard was required.

3.5.5 Contact form

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

Besides 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 (see online appendix) 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 advice on how to find the addresses of households and convince them to take part in the HFCS. It likewise described the requirements for interviewer behavior and their interaction with the people contacted. Other important features were detailed contact provisions and answers to questions frequently asked during the first contact attempt. The manual additionally comprised essential legal texts regarding the guarantee of 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 the interviewers' 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 households were forwarded to the OeNB promptly, in 15 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) based on household sheets which contained household level information from interviews and were evaluated individually and grouped by interviewer. 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 duration¹ 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.

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

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 a few cases, misunderstandings were found in relation to the reference period for income (the gross income recorded referred to the 2013 calendar year for most of the surveyed income variables) or with regard to other flow variables (e.g. the variable for current consumption expenditure, where the reference period was a typical month). In the case of some of the interviewers, monitoring also helped reduce the item non-response rate of the households they interviewed.

Four interviewers had to be withdrawn entirely from the survey during the fieldwork due to flaws in conducting the interviews. In these cases, the household data potentially containing quality flaws were subjected to increased scrutiny. Missing information was collected through follow-up queries; if there was still considerable doubt concerning the quality of the information collected, the interviews in question were classified as unsuccessfully completed after the field phase. A total of 38 households that had participated in interviews were excluded from the data due to qualitative flaws.

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 volun-

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

tary 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). 55 of the 72 interviewers took part in the survey of interviewers.

3.9 Online appendix

The online appendix includes the letter by the OeNB governor to the households, the information leaflet, the catalogue of showcards, the interviewer manual, the glossary, as well as the exemplary study by Fessler et al. (2012).

4 Consistency checks and editing

4.1 Introduction

Data editing is understood to mean the later 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.

This chapter provides insights into the consistency analyses and edits performed for the second 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 65,000 of the close to 1.3 million observations collected in the second HFCS wave were edited, i.e. 4.8% of all data points are amended (see table 2).

Table 2

Number and type of edits

	Total observations ¹	Number of edits	Share of edited observations in total observations
All	1,347,529	65,237	4.8%
Edits based on expert judgment and follow-up phone calls	1,347,529	9,319	0.7%
Edits based on other survey information (e.g. verbatim records)	1,347,529	44,715	3.3%
Deleted observations	1,347,529	11,203	0.8%

Source: HFCS Austria 2014, OeNB.

¹ Includes only observable information. Filter missings are excluded.

The rows below “All” indicate the different types of edits. Edits resulting in changes to the collected values, i.e. real changes, occurred in the case of only

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.

around 9,300 observations (see row “Edits based on expert judgment and follow-up phone calls”), which corresponds to a change rate of 0.7%. These changes involved primarily inconsistent values that were corrected as a result of subsequent queries on the phone and/or other information, or were deleted and replaced through imputation. More than two-thirds of all amendments (see row “Edits based on other survey information (e.g. verbatim records)”, i.e. about 44,700 observations, could be derived from the verbatim records and the use of respondent-friendly questionnaire design (e.g. questions about life insurance policies or total annual net income). All in all, some 3.3 % of all observations were amended through this type of editing. This rate 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. In around 11,200 cases (i.e. around 0.8% of observations), observations were set to filter missing (“.”),² mostly in the process of data cleanup (see section 4.6.2.1). Moreover, some items had been entered at a wrong position in the questionnaire. When transferring such information to the right position, the original entry must be deleted. In addition, there are cases (see below) in which a complete observation is set to filter missing (“.”) among other things because the corresponding head variable has been edited.

A case in point³ 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. The interviewers use a laptop on which the survey software is installed and are guided through the questionnaire on the screen. 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. Interviews might even have to be broken off in individual cases.

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 indented for ease of reference.

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 the latter 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. Whenever a test criterion was violated, a warning message popped up.

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, these exemplary 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).

4.4 Postinterview consistency checks

4.4.1 Expert data analysis

During the field phase of the second HFCS wave in Austria, the data of households deemed to be final by the survey company were forwarded to the OeNB in 15 batches. This means that the OeNB received household data roughly every three weeks during fieldwork. All batches of data were subjected promptly to

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.

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. Thus 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. housewife or pensioner bias) and distorted the data accordingly. The interviewers knew that the list of addresses was limited to the 6,308 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. 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 geographical location and structure, financial and real assets, debt and income, whether households had come to 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 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 four interviewers whose results were not up to the required standards (e.g. regarding nonresponse) were excluded.

4.4.2 Follow-up queries

If individual data analysis did not reveal the type of problem or 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 three weeks) and subsequent check by the HFCS team, the survey company was able to address any queries to the surveyed households promptly. A typical case of a data problem that

⁵ The data evaluation was conducted with the aid of the results of the first HFCS wave as well as external data sources such as the EU-SILC (conducted by Statistics Austria).

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 follow-up telephone 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 else corrected in the course of queries. All in all, follow-up queries (by phone) were necessary to confirm specific details of some 400 households, which is a smaller percentage of households than in the first wave. This decrease is, above all, attributable to the substantial increase in the use of comment fields (as a result of the experience gained in the first 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). These outliers were recorded above all for wealth variables, the size 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 largely 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.

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 55 households made it possible to largely exclude programming errors from the outset. Minor difficulties, e.g. incomplete filtering with regard to the question on additional borrowing (HB150\$x) (see section 2.5.2.3) were 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). To comply with interna-

⁶ This problem resulted from an update of the questionnaire during the field phase and was relatively difficult to identify as previous interviews showed correct filtering techniques before the update.

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).

tional requirements, some flags were aggregated for the international datasets (section 4.7). The flags used can be divided into five groups.

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 filter missing values (“.”) were flagged “0.” Information recorded in loops (see section 4.6.2.4) was paired with a flag of 2 if it had to 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. The yes/no question on silent partnerships (HD1000) was encoded with “yes” for those households that held investments in a self-employment business, but had no active role in running the business and were not self-employed in this business. Those (few) observations in this variable were given a flag of 12.

Table 3

Flags used in the HFCS in 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 the 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 2014, OeNB.

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 observed value. 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 information from telephone 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.

Chart 3 indicates how questions were typically structured in the HFCS questionnaire. Let us take employee income to give an example for the structure of question blocks⁸ and the use of flags.

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 in the next question 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 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

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

Sequence of questions



Source: HFCS Austria 2014, OeNB.

response to a subsequent question (e.g. on employment) revealed ex post that a “No” given to 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 variable for the value of the income was flagged for imputation. Following imputation, the value was then reflagged “4056” (Imputed, originally 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 every single observation in the HFCS to be tracked. 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 that 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 judgment was 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, the changes were made.
- Data entry errors by interviewers: In one instance, the contact month for an interview conducted in 2014 had been entered as January 2014 (“1”). This was changed to October 2014 (“10”), because the data on the relevant household were submitted in November and the preceding and following contact attempts had taken place in October 2014.
- 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 when they proceeded backward 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 cleaned up (i. e. observations needed to be deleted ex post).

4.6.2.2 Currency conversion

Respondents could specify any amount in any currency (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 1.95583.⁹

In a few cases – in particular for foreign currency loans – amounts were also given in Japanese yen and Swiss francs. The value of the amount outstanding at the time of the interview was converted into euro on the basis of the average 2014 exchange rate, while the total value at the time of borrowing was converted at the average of the exchange rates recorded in the year in which the loan was taken out, with the exchange rates published on the OeNB’s website¹⁰ being used as exchange rates.

Individual cases¹¹ involving other currencies were converted with great care using the annual average of the respective currency’s exchange rate to the euro. Two inheritance cases from before the introduction of the euro were first converted from Deutsche mark into Austrian schillings on the basis of the applicable exchange rate at the time and then from Austrian schillings into euro according to the fixed ATS/EUR exchange rate.¹² Likewise early amounts in terms of the reference year in Canadian dollars were first converted into Austrian schillings and then into euro.¹³

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:

⁹ See <https://www.oenb.at/isaweb/report.do?jsessionid=31767F3B9E6FA661A8A4CD5CB700B5A7?report=2>.¹² (accessed on December 9, 2016).

¹⁰ See www.oenb.at/en/Statistics/Standardized-Tables/interest-rates-and-exchange-rates/Exchange-Rates.html (accessed on December 9, 2016).

¹¹ Some observations of variables specific to Austria that are not contained in the internationally available core dataset required the use of historical OeNB exchange rates. However, these are not included in this documentation, as they are not part of the available data.

¹² The relevant exchange rate was taken from the OeNB’s *Statistisches Monatsheft* of December 1998 (OeNB, 1998).

¹³ The relevant exchange rate was taken from the *Mitteilungen des Direktoriums der Oesterreichischen Nationalbank* 1979 (OeNB, 1979).

- 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
- 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 course of the editing process – with the exception of the loop questions on inheritances, for which no recoding was carried out because respondents were prompted to record the inheritances received in descending order of relevance for the household’s current wealth situation. They were, however, 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 filter missing in one iteration was replaced with the same variable set to filter missing in another iteration, it was flagged “0” (“Not applicable (skipped due to routing)”).

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 and private businesses, the relevant summary responses had to be edited accordingly. For ease of reference, examples of these edits are described below on the basis of the section of the questionnaire dealing with other unsecured loans (see section 2.5).

In the 18 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 answer to the summary question. If the respondent had indicated the outstanding amount only in their answer 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.¹⁴ 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,¹⁵ 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. 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. This case did not occur in the example given above.

All edits were again flagged correspondingly.

¹⁴ 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.

¹⁵ In the second HFCS wave in Austria, only one household opted for this route.

Summary questions

Every loop of questions ended with summary questions (see chart 2 in chapter 2). The variables for these questions exclusively contained information on any additional items above three per household. As indicated in chart 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 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 (see also section 2.6.3) proved very helpful during the various steps of data processing.

Some 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.” Use of the additional guidance to the respondents during the interview made it possible to reduce the number of such cases considerably compared with the first wave.

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 8 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 the list of individual IDs that were required for the internationally available dataset (which only contains six variables for individuals).

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 not very error-prone. In particular, there was 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 cumulative sum of all payments. In cases where one or several details were not given, the remaining observations were used as bounds for the value to be imputed. Insurance policies (term-life insurance) 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 old 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 real estate assets (HG0300 and HG0310)
- income from financial investments (HG0400 and HG0410)
- income from private 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,¹⁶ 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 as a blue-collar or white-collar worker or retired).¹⁷ 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 recorded only for annual incomes less than EUR 11,000, which are classified as tax-free, so that the gross amount is equal to the net. For all other values (for some 45 individuals), a range was created for imputing specific amounts by adding EUR 10,000 to and by subtracting EUR 10,000 from the amount converted subject to the conditions for white-collar workers. 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. In all conversions, all types of income were always included together. If a respondent indicated both employee and self-employed income, the sum of both incomes was used and converted. The employment status of the respondents was dependent on their main source of income. The total gross value was then split in proportion to the net income proportions indicated. To calculate the gross income from financial investments, 25% withholding tax (capital income tax) was added to amounts given for net income.

¹⁶ See www.bmf.gv.at/service/anwend/steuerberech/bruttonetto/_start.htm (German only) (accessed on December 9, 2016).

¹⁷ "Apprentices" were categorized as "blue-collar workers" in the conversion, while "civil servants" were seen as "white-collar workers" on grounds of their more favorable tax treatment.

Table 4

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

	Number of persons	Share in %
Number of persons receiving employee income	2,638	100
Answer recorded, complete observation (flag 1)	1,067	40.4
Not imputed, originally "Don't know" (flag 1050)	32	1.2
Not imputed, originally "No answer" (flag 1051)	105	4.0
Not imputed, originally collected from a range (flag 1053)	326	12.4
Not imputed, collected value deleted (flag 1054)	7	0.3
Not imputed, set to missing due to incorrect answer to a higher-order question (flag 1056)	36	1.4
Not imputed, collected value deleted but range information available (flag 1057)	66	2.5
Edited, set to modified value as considered incorrect or unreliable (flag 3050)	1	0.0
Edited, adjusted on the basis of other collected (national) variables (flag 3051)	965	36.6
Edited, adjusted on the basis of the verbatim records (flag 3052)	2	0.1
Edited, adjusted on the basis of follow-up with household (flag 3075)	30	1.1
Edited, adjusted on the basis of follow-up with interviewer (flag 3076)	1	0.0

Source: HFCS Austria 2014, OeNB.

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 2,638 individuals. 1,067 respondents (40.4%) expressed their annual income in gross terms. A further 32 respondents (1.2%) answered "Don't know" and 105 individuals (4.0%) opted for "No answer." 326 respondents (12.4%) specified their income amount using a range. For 66 individuals, a range could be calculated from other information given. The responses of 43 individuals (around 1.6%)¹⁸ were edited, set to missing and flagged for imputation; the vast majority of these edits (for 36 individuals) were due to an incorrect head variable (flag 1056). 36.6% of the respondents (965) provided their net income, which was then converted with the aid of the finance ministry's gross-to-net converter. Expert judgment was used to edit the income of one individual. The responses of the remaining 33 individuals (around 1.2%) were corrected on the basis of follow-up queries and verbatim records.

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

¹⁸ Different from table 4 on account of rounding differences.

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 codes, using the German version of Statistics Austria's information on ISCO-08.¹⁹ 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 the ISCO classification (in particular, the respondent's level of education and the main activity of the company where the respondent worked). The variable PE0300 to be submitted to the ECB was first flagged "3052" (Edited, adjusted on the basis of the verbatim records 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.²⁰

4.6.2.12 Highest education qualification

To account for latest developments in the *International Standard Classification of Education* (ISCED), the highest education qualification of all household members was recorded in substantially more detail than during the first HFCS wave. Respondents were not asked about ISCED categories, however, but were prompted to indicate their qualification based on Austria's education system. Additionally, the bachelor, master (including Magister and diploma degrees) and doctorate were recorded. The bachelor degree may have been inadequately chosen by some respondents. As this is a fairly new degree in Austria, respondents aged 40 or above are rather unlikely to have graduated with this degree. Therefore, these individuals were assumed to have a master degree (including Magister and diploma degrees); the corresponding variable was flagged "3051." As the national degree hierarchy was aligned with the ISCED codes (see also section 4.7), this aspect does not play a role for the international dataset, though.

4.6.2.13 Exclusion of successful interviews

For various reasons, the final data do not include those from altogether 42 households that were interviewed (successfully).

- Households not belonging to the target population: The target population for the HFCS in Austria comprises all households that do not live in institutions (e.g. children's homes, retirement homes, prisons). A number of households that had been interviewed successfully were excluded from the survey because the respondents were living in student housing. The four households in this group were edited out and flagged as "Not belonging to the gross sample." 33 other households that were not interviewed successfully were likewise eliminated from the gross sample. In addition to households living in institutions, these households also comprised addresses whose occupants had died and addresses that were used commercially.

¹⁹ For further information, see www.statistik.at/web_de/klassifikationen/oeisco08_implementation/informationen_zur_isco08/index.html (German only) (accessed on December 9, 2016).

²⁰ For explanations, see www.statistik.at/web_en/classifications/implementation_of_the_onace2008/index.html (accessed on December 9, 2016).

- Households with an excessive proportion of nonresponse items: This group, which comprised 38 households, had to be deleted from the dataset because the respondents refused to answer too many questions. The observations for these households were edited to “Interview completed, but rejected after fieldwork” and assigned a nonresponse weight of zero (see section 7.2.3).

4.6.2.14 CAPI errors encountered with the questionnaires

Additional borrowing (HB150\$x)

Because of a CAPI error, a total of 39 households were not asked the question on additional borrowing (HB150\$x), although they should have been routed to this question. These observations were set to missing with a flag of 1055 and released for imputation. Following imputation, they were reflagged “4055.”

The few other problems encountered with the programming of the questionnaire (section 2.5.2.3) did not require any further editing measures.

4.7 Formatting and editing after multiple imputations

Any information collected at a greater degree of granularity in Austria was processed further upon imputation so as to bring the level of aggregation into 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.”
- Education: Categories specific to Austria were paired with ISCED 1997 codes²¹ and classified as ISCED level 0 (“Compulsory education not (yet) completed”); ISCED level 2 (“Compulsory education completed”); ISCED level 3 (“Apprenticeship (vocational school)” and “Other vocational middle school”); ISCED level 4 (“Nurses’ training school”, “Secondary academic school – senior classes” and “Vocational or technical school – college, high-school graduates’ training course”); ISCED level 5 (“Master craftsman – works master training” as well as “University, academy, technical college – bachelor” and “University, academy, technical college – Magister/diploma degree/master”) and finally ISCED level 6 (“University: doctorate”).
- 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.
- Loan repayments: The installments for repaying (secured and unsecured) bullet loans were set to “0” as such loans are repaid with a single lump sum upon maturity. Assets accumulated for repayment can be analyzed on the basis of variables that are specific to Austria.
- 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.”

²¹ For explanations, see www.statistik.at/web_de/klassifikationen/klassifikationsdatenbank/weitere_klassifikationen/bildungsklassifikation/104092.html (German only) (accessed on December 9, 2016).

- Number of other vehicles: The vehicle categories “Vans” and “Mobile homes and caravans” were aggregated as “Vans.”
- Purpose of a loan: The category “To finance a deposit for the housing association” was 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.
- Silent partnerships: The yes/no question recording silent partnerships was not put to households that held investments in a self-employed business but were not actively involved in its management or working for the business on a self-employed basis. For reasons of consistency, the respective variable (HD1000) was encoded with “1 – yes” ex post and flagged “12.”
- Investments in savings plans with building and loan associations and life insurance policies: Data recorded on these two investment methods were aggregated into savings (HD1200 and HD1210).
- 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.
- Purpose of saving: The sequence ordered by relevance was abandoned.
- Paradata: The variable on alarm systems and other security measures (ASC0700a-h)²² was aggregated (the possibility of multiple responses was abandoned).

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 1057 was recoded as 1053
- flag 1058 was recoded as 1051
- flags 3051, 3075, 3076, 1075 and 2 were recoded as 1
- flag 3052 was recoded as 11
- flag 3053 was recoded as 0
- flag 4057 was recoded as 4053
- flag 4058 was recoded as 4051

The following variables are exceptions to this rule:

- Income variables²³ flagged “3051” after net-to-gross conversions were reflagged “13” (rather than “1”) following aggregation in line with international requirements.
- The variable featuring rent excluding utilities (HB2300) was flagged “1075” to identify the special cases when only the rent including running costs 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 wave, 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 spring 2017. The additional information includes additional variables, as well as a

²² These variables are not contained in the international core dataset.

²³ This concerns the following variables: PG0110, PG0210, PG0310, PG0410, HG0410.

more detailed breakdown of certain variables. Datasets may be merged on the basis of both the identification numbers and imputation numbers.

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 follow-up inquiries 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 the analysis of 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.²⁴

²⁴ *All documents included in the online appendix are available at www.hfcs.at/en.*

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).

If the problem of missing information due to item nonresponse were disregarded, it 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).

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) – also use the same approach to impute missing data.

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 specification of the imputation model 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. On average, each household has 29.9 missing values, which means that item nonresponse was limited to 2.1% of all the questions (variables) addressed to each household. However, the respective percentage for the euro variables amounts to 4.7%. This suggests that questions of this kind might be perceived as sensitive or difficult to answer.

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 on complete

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 (2002).

Table 5

Item nonresponse per household (unweighted)

	Mean	Median	Minimum	Maximum
Number of variables asked				
all variables	1,392.0	1,391.0	1,109	1,889
euro variables	63.0	64.0	36	106
Number of variables with missing values				
all variables	29.9	18.0	0	487
euro variables	3.0	2.0	0	36
Share of variables with missing values in %				
all variables	2.1	1.3	0.0	32.0
euro variables	4.7	3.0	0.0	49.2

Source: HFCS Austria 2014, 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.

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.

Table 6 shows that for example when asked about the value of their main residence, 77.1% of households provided a specific amount (column 3). The other

Table 6

Item nonresponse for selected variables (unweighted)

	Household has item		Responses by households that have the item			
	Yes	Unknown	Amount	Range	"Don't know"/ "No answer" (5)	Other missing values ¹ (6)
	(1)	(2)	(3)	(4)	(5)	(6)
	%					
Value of main residence ²	42.9	0.0	77.1	19.4	3.2	0.4
HMR mortgage 1: amount still owed	12.7	0.6	69.4	14.9	15.4	0.3
Monthly amount paid as rent	50.9	0.0	61.7	37.8	0.4	0.1
Other property 1: current value	11.1	0.3	78.4	14.7	6.0	0.9
Other property mortgage 1: amount still owed	1.3	0.3	77.5	10.0	12.5	0.0
Value of sight accounts	99.2	0.0	80.4	9.5	9.5	0.6
Value of saving accounts	84.0	1.2	72.5	14.2	9.6	3.7
Value of publicly traded shares	5.0	0.4	66.2	13.3	18.5	2.0
Amount owed to household	7.6	0.4	94.3	2.6	3.1	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)	48.7	0.0	85.0	9.9	3.7	1.4
Gross income from unemployment benefits (person 1)	5.5	0.1	89.8	5.4	3.0	1.8
Gross income from financial investments	63.0	15.0	46.6	33.0	19.1	1.3
Gift/inheritance 1: value	26.7	1.3	77.8	9.6	9.0	3.5
Amount spent on food at home	100.0	0.0	98.4	1.4	0.2	0.0

Source: HFCS Austria 2014, OeNB.

¹ Missing values due to editing measures and exits from loops.

² Based on the HB0900 variable.

Note: HMR = household main residence.

22.9% of households are item nonrespondents: Either they provided a (prespecified or individual) interval (19.4%, column 4), responded with “Don’t know” or “No answer” (3.2%, column 5) or their response was set to missing (0.4%, column 6).³ Nonresponse rates⁴ vary substantially across items. Variables with high nonresponse rates include e.g. questions related to the value of publicly traded shares ($100\% - 66.2\% = 33.8\%$) and the household’s gross income from financial investments ($100\% - 46.6\% = 53.4\%$). With regard to the latter, 33% of households provided at least a range for this type of income, 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). A variable with a low nonresponse rate is, e.g., the amount spent on food consumed at home ($100\% - 98.4\% = 1.6\%$).

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 (63%) are then asked the question on the amount of income; the other households, including the 15% of households that did not answer the yes/no question, automatically skip the euro question. As it is unknown, however, whether the 15% 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.

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 also because of 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 illustration purposes, table 7 shows a regression of nonresponse to

Table 7

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

Covariates	Coefficient
Female (person 1)	-0.0777 (0.0966)
Age (person 1)	-0.00382 (0.00338)
Tertiary education level (person 1)	-0.00624 (0.127)
Employed/self-employed (person 1)	-0.315*** (0.115)
Residence is in Vienna	-0.642*** (0.132)
Size of main residence	0.00225** (0.000987)
Household size	0.282*** (0.0440)
Constant	-1.750*** (0.248)
Observations ¹	2,940

Source: HFCS Austria 2014, OeNB.

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

¹ The remaining 57 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.

³ See chapter 4 for more details.

⁴ The nonresponse rate is calculated by subtracting the value in the “amount” column in table 6 from 100%.

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 live in smaller main residences and in smaller households, they tend to live in Vienna and are more often gainfully employed. Thus, a complete-case analysis of the value of sight accounts would bias the estimates toward a population with these household characteristics.

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 imputed 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, \dots, Y_P to be imputed.
- Step 2: Fill the missing values of Y_1, Y_2, \dots, Y_P with random selected values which were actually observed.
- Step 3: For each Y_1, Y_2, \dots, 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 distribu-

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, which is identical to that used for the HFCS.

tion of variables with missing values. If t tends to infinity, the imputations of missing values of Y_1, Y_2, \dots, 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, \dots, Y_p converge in expectation 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, filter missings, bracketed (i.e. range) responses, interactions or constraints in relation 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, \dots, Y_p to be imputed. Our strategy is to impute as many variables with missing values as possible, which amounts to around 70% of such variables. The remaining variables with missing values are not imputed with the HFCS imputation procedure due to a lack of sufficient variance or due to 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

⁷ See Little and Rubin (2002) 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.

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 section 5.3 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, \dots, Y_p often aids convergence. Therefore, we order the variables to be imputed by their degree of missingness, starting with the variables with the least 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 was always imputed before the mortgage amount was imputed, even if the degree of missingness was the same for both variables.

5.4.3 Types of regression models

In step 3, we defined 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 used an interval 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 used a logit model; and for ordinal and nominal variables, we used 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. In the second wave, we decided to use weights as predictors (section 5.4.7), but not for weighted regressions, as is the current practice in imputation (see e.g. Frumento et al., 2012). The reason stated in the literature is that 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.

⁹ The interval 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 to first, predict a value by linear regression for each missing value and second, impute the observed value that is closest to the regression-predicted value.

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 is the result of using the natural logarithm. These types of variables usually have a highly skewed distribution; using the logarithm helps to make the distribution closer to the normal distribution assumption that is necessary for the prediction. Another very helpful transformation for year variables is to impute time periods instead of 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. 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, 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 interval 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 people 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 already 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 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 interval 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.

General bounds are the same for all households and persons, and 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 in Austria was taken out, renegotiated or refinanced more than 70 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 for part-owners), where we set the lower bound to zero and the upper bound to 100, or some year variables (e.g. the purchase year of the household's main residence), where the upper bound is 2015, i.e. the year in which the last survey interviews were carried out.

¹¹ In the dataset, financially knowledgeable persons are designated with the ID = 1 by default; all other people are ordered by age.

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 on which our imputation model relies (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 instance, variables that describe the household (e.g. estimated household income, household size, number of children), variables that describe household members (e.g. age, education, gender and employment status of the household's first individual and his/her partner), stratification variables (e.g. province, municipality size), information provided by the interviewers (e.g. standard of living, type of neighborhood, building condition, interview atmosphere, etc.). The latter pieces of information (para-

- data) 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, the length of time (in years) for which the household has already owned the respective property and the total value of real estate property owned by the household. Usually, these variables are connected logically (e.g. 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 different 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 above-mentioned predictor variables that households that were asked about the variable to be imputed were not asked about. 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 mort-

gage amounts. If, for those households that do not have a mortgage, we just set the mortgage amount to zero (which 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), and which 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 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 cycles. Given the large computational effort required for the HFCS imputation model, we set the cycle number for

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.

the HFCS imputation model at $t=10$. Other similar surveys, like the SCF (Kennickell, 1998) and the EFF (Barceló, 2006) use $t=6$.

Typically, we check convergence graphically by plotting the mean of the imputed values against the cycle 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.

In the second wave of the HFCS, we additionally examined the convergence of selected variables using the Gelman-Rubin convergence diagnostic, which is used very frequently in 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 between the multiple imputation samples compared to the variance of the same estimate between the cycles.¹³ All variables examined in the second wave of the HFCS meet this criterion.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 $m = 1, 2, \dots, M$ 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 claim that good inferences can already be made with $M=3$ to $M=5$. In line with the international requirements and 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.

One interesting result is that the means of most variables 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 87,202 before imputation. After the respective imputations, it

¹³ The Gelman-Rubin diagnostic is the 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 to 1.1, they are usually considered to denote convergence.

¹⁴ The following important variables were tested: HB0900, HB1701, HB2801, HB4400, HD1110, HD1210, HD1510, HI0100, HI0200 and HI0310.

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	285,996	290,833	290,995	292,706	290,210	292,890
HMR mortgage 1: amount still owed	73,205	80,468	86,705	81,603	81,151	85,650
Monthly amount paid as rent	407	393	399	401	396	391
Other property 1: current value	249,384	237,947	248,696	258,458	233,246	246,517
Other property mortgage 1: amount still owed	78,480	81,357	70,089	74,470	67,089	74,713
Value of sight accounts	2,623	2,689	2,695	2,624	2,612	2,528
Value of saving accounts	23,201	26,925	27,293	26,375	26,526	27,389
Value of publicly traded shares	27,584	26,222	32,490	25,038	26,693	31,007
Gross cash employee income (person 1)	27,319	27,677	27,587	27,695	27,509	27,560
Gross income from unemployment benefits (person 1)	6,437	6,504	6,502	6,363	6,482	6,664
Gross income from financial investments	706	523	564	553	596	587
Gift/inheritance 1: value	87,202	92,620	91,502	92,076	100,621	97,088
Amount spent on food at home	373	374	373	373	373	373

Source: HFCS Austria 2014, 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.

increases to EUR 92,620 in $m = 1$, EUR 91,502 in $m = 2$, EUR 92,076 in $m = 3$, EUR 100,621 in $m = 4$, and EUR 97,088 in $m = 5$. Thus, on average the imputations increase the mean value of the first gift/inheritance from EUR 87,202 to EUR 94,781, i.e. by 9%. Additionally about one-third of the values imputed in this context are based on interval responses by households, which suggests that households with more valuable inheritances tend to answer with an interval response to this question less often than households with less valuable inheritances. The largest increases in comparison to the unimputed sample occur when imputing savings account holdings and mortgage loans for the main residence. Households' interval responses again play an important part here, as they provide valuable and often very precise information for the imputations (see also table 6).

However, for some variables, the mean does not change significantly and, in some cases, it even decreases. 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 that nonrespondents with regard to this variable tend to have lower income from financial investments.

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 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.

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[®], please see the HFCS User guide (chapter 9).

6 Sampling

6.1 Introduction

The sampling design for the second wave of the HFCS in Austria was specifically developed by the OeNB in collaboration with the survey company IFES (Institut für empirische Sozialforschung GmbH). 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. The units of the sample should be representative of the whole population; in other words, in expectation an analysis of the sample (using appropriate weights) should 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 regions – is achieved through stratified sampling, i.e. by dividing the country into smaller geographical 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 the best unbiased estimates (and confidence intervals) and keep uncertainty as low as possible by taking into account restrictions like costs, time and practicability. 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 a box. 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 defining the target population of the survey. The HFCS is intended to cover all households living permanently in Austria, independent of citizenship and/or residence status. According to the ECB, 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, the following persons are to be regarded as household members, according to the ECB's definition¹:

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

¹ See ECB (2013a), p. 80f.

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”

In the case of the HFCS, the target population does not include households that live in institutions such as:

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

Additionally, the Austrian HFCS does not cover homeless people. People without a residence could not be reached with the survey, as sampling was based on 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.

Box 1


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 units of collection – i.e. households, for our purposes – are drawn from all subgroups of the target population. Stratification in the Austrian HFCS was carried out geographically (based on NUTS-3 regions¹) and by municipality size categories.

“Two-stage cluster sampling” means that, first, primary sampling units (PSUs) are selected from each geographical unit (i.e. stratum) and that, second, secondary sampling units (SSUs) are 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. Unlike during the first wave of the HFCS in Austria, the probability of being drawn during the first stage is proportional to the number of households in the respective PSUs. 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.

¹ See www.statistik.gv.at/web_en/classifications/regional_breakdown/nuts_units/index.html (accessed on December 9, 2016). 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.



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 HFCS in Austria 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 the 2011 register-based census.² 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 440 dwellings on average).³

² See also www.statistik.at/web_de/statistiken/menschen_und_gesellschaft/bevoelkerung/volkszaehlungen_registerzaehlungen_abgestimmte_erwerbsstatistik/index.html (German only) (accessed on December 9, 2016).

³ The estimated number of households in Austria according to the HFCS definition (3.9 million) divided by the number of enumeration districts (8,821) yields 442.1.

In addition, we relied on the municipality directory of 2013 (to categorize by municipality size) and on Statistics Austria's register-based census from 2011 for the population data.⁴ Both datasets were collected and provided by Statistics Austria. The register-based census of 2011 indicates in particular the population (households) of each stratum and thus serves to determine the number of sampling units to be drawn within each stratum. To sum it up, all the information necessary for the stratification and the first stage of the sampling design is based on geographical and population information provided by Statistics Austria.

6.3.2 Austrian Post Office

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 Office, 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 Office. 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 which mail can be delivered to. This information is available in disaggregated form. In Austria, there are about 4,051,000 private mail delivery points. Addresses identified by the Austrian Post Office as vacation homes have already been excluded.

Thus, our starting point was some 4.1 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 arrived at the address and noted that it was incorrect or was 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 3.86 million; thus, Austria has an estimated 3.9 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.⁵

⁴ See also Statistics Austria (2013).

⁵ The post-certified addresses for some 4.1 million households compare with about 3.8 million household addresses documented by other sources (such as the microcensus based on the residence registry; see also section 7.2.4).

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.⁶

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 chart 4).


With the exception of the capital city Vienna, each NUTS-3 region was divided further into the following eight categories according to municipality size.

- 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 national capital and the only city with more than 1 million inhabitants; it was subdivided into its 23 districts.

Chart 4

NUTS-3 regions



Source: Statistics Austria.

⁶ For more details on the contact strategy, see section 3.4.

This very fine stratification yielded 193 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 with neighboring strata to increase the share of households and thus insure the selection of at least one PSU from each stratum. This exercise left the HFCS with 185 strata for sampling covering all households in Austria. The distribution of strata across provinces and municipality size categories can be seen in table 9.

Table 9

Allocation of strata within the sample

	Municipality size ¹ (number of inhabitants)								Total
	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	
Vienna	0	0	0	0	0	0	0	23	23
Lower Austria	7	7	7	7	6	5	1	0	40
Burgenland	3	2	2	2	1	0	0	0	10
Styria	6	6	5	6	2	1	1	0	27
Carinthia	3	3	3	3	2	1	1	0	16
Upper Austria	5	5	5	5	3	2	1	0	26
Salzburg	3	2	2	2	2	1	1	0	13
Tyrol	5	5	3	3	3	0	1	0	20
Vorarlberg	2	2	2	1	2	1	0	0	10
Total	34	32	29	29	21	11	6	23	185

Source: Statistics Austria (municipality directory 2013).

¹ Municipality size accounts for municipality mergers up to 2013.

Each stratum contained about 50 PSUs on average, which in turn contained around 460 households⁷ on average.

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 first wave of the HFCS in Austria in 2010 and 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).⁸ With some leeway for extraordinary circum-

⁷ The difference between this figure and the 440 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 3.9 million households according to the HFCS definition in the 8,471 enumeration districts resulting from the aggregation to calculate this average.

⁸ The regional differences between participation rates in the first wave of the HFCS in Austria were taken into account in the sample design for the second wave.

stances, the HFCS survey was designed to yield a sample of 3,000 successfully interviewed households and a participation rate of about 35% in Vienna and around 55% in the rest of Austria. These participation rates are estimates based on the experience of past surveys. The participation rates recorded in the first wave in a stratum were used to determine the exact number of households to be drawn in the respective stratum in the second wave.

The targeted net sample of $n = 3,000$ was divided between the nine provinces, based on their population share as surveyed by Statistics Austria in the register-based census of 2011⁹ (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).

Table 10

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

	% of households (1)	Target sample (2)	Gross sample (3)	Number of households per PSU (enumeration district) (4)	Number of PSUs to be drawn (5)
Vienna	23	688	1,960	8	245
Lower Austria	19	599	956	8/12	82
Burgenland	3	96	168	12	14
Styria	14	421	832	8/12	82
Carinthia	7	198	404	8/12	41
Upper Austria	16	485	932	8/12	90
Salzburg	6	184	384	8/12	39
Tyrol	8	244	468	8/12	45
Vorarlberg	4	125	204	12	17
Total	100	3,000	6,308	29	655

Source: Statistics Austria (municipality directory 2013), HFCS Austria 2014, OeNB.

In total the Austrian HFCS sample design produced 655 (619 unique) PSUs across all strata and a gross sample size of 6,308 households that were invited to participate in the HFCS (see box 2 in chapter 7 for information on the number of households interviewed successfully). Drawing the PSUs was done with replacement, which entails some PSUs being drawn multiple times (see section 6.5.1). Drawing possible substitute addresses was explicitly 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).

⁹ Although Statistics Austria's register-based census of 2011 is based on a different definition of households, we used the respective data as an approximation.

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 8,821) as the PSUs for the Austrian HFCS. On average, an enumeration district contains 440 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 8,471, which then had 460 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 90 PSUs to be drawn in Upper Austria (table 10) were divided up among the 26 strata in this province according to the population share.

After determining how many PSUs were to be drawn per stratum, the PSUs – unlike during the first wave of the HFCS in Austria¹⁰ – were drawn proportionally to their size (measured in terms of the number of households in a PSU).¹¹ The purpose of this change is to reduce the standard errors of estimates by reducing the variance of the sample design weights (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 655 PSUs were drawn in the second wave of the HFCS in Austria, only 619 of which were unique.

6.5.2 Second stage

With 655 (619 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 6,308 households in Austria.

¹⁰ During the first wave, the probability of drawing a PSU within a stratum was identical for every PSU.

¹¹ 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.

6.5.3 Practical implementation

Table 11 illustrates how the data from the second stage of the sampling were used, after the PSUs had been chosen in the first stage (column 2): Austrian Post Office 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 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).

Table 11

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

First stage		Second stage				
Statistics Austria		Austrian Post Office				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 Office, Profile.Address/IFES.

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).¹²

6.6 Concluding remarks

This chapter provides information on the sampling design specifically developed for the second wave of the HFCS in Austria on the basis of the first-wave design. The improvements specified above were implemented as a result of the experience and information gained from the first wave. 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 Office database) from the selected PSUs.

The sampling method used for the HFCS has a number of advantages, with the following aspects being particularly important:

- Compared to the first wave, the fact that the probability of being drawn was proportional to the size of the PSU in terms of the number of households improved the efficiency of sampling design by reducing the variance of the design weights.

¹² See the online appendix for the invitation letter.

- 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, the underlying information needed for such oversampling is not available yet.

7 Construction of survey weights


7.1 Introduction

Survey weights are usually computed 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 chart 5).

Chart 5

Misrepresentation of the target population in the sample



Source: Adapted from Biemer and Christ (2008).

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 that the number of primary sampling units (PSUs) to be drawn per stratum is fixed by the HFCS sampling design. Another example is the over-sampling of households in metropolitan 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 computed 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

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.

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, for example, those of companies² or individuals living in care residences. 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 weights³ (to address erroneous inclusion and frame multiplicity) or poststratification weights (to address erroneous exclusion). We explain the computation 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 chapter 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 computed and how the final set of survey weights was derived from these weights. Moreover, 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 misrepresent the target population for various reasons. Therefore a specific adjustment using weights

² Although the sampling frame was cleaned of addresses of companies, some may still be erroneously included.

³ Sometimes referred to as noncoverage weights.

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.

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 :

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

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 in 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 section 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 an 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 every household i is given by multiplying the two partial probabilities. This probability may be shown as:

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

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 4.1 million to 3.9 million.

Table 12

HFCS design weights by federal province

	Mean	Median	Minimum	Maximum
Vienna	494	513	0	707
Lower Austria	757	779	0	1,099
Burgenland	729	694	0	1,008
Styria	641	636	0	1,146
Carinthia	613	484	0	1,183
Upper Austria	626	632	0	1,004
Salzburg	600	512	0	925
Tyrol	651	624	0	1,085
Vorarlberg	792	833	0	1,045
Total	614	547	0	1,183

Source: HFCS Austria 2014, OeNB.

Table 12 shows some statistics of the obtained HFCS design weights across Austria's provinces. Vienna and Salzburg are the provinces with the lowest median weights, which is plausible, as households living in these regions were oversampled because of their low willingness to participate during the first HFCS wave in Austria, which would have created a bias had they not been reweighted downward using the design weights.

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 513 households in the sampling frame.

Box 2

Unit nonresponse in the HFCS in Austria

In the HFCS in Austria, successful interviews were conducted with 2,997 households from the gross sample, which comprised 6,308 addresses. The remaining 3,311 addresses were classified either as unit nonresponse (2,997 households), ineligible addresses (284 addresses) or addresses of unknown eligibility (30 addresses).

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 2,657 households. Another reason was that no contact at all could be established with 136 households. The remaining 204 nonrespondents specified other reasons, such as illness, language barriers; or they resulted from ex post exclusion of interviews due to a high number of missing or unreliable values.

In addition, 284 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 another 30 addresses was impossible to ascertain, as the interviewers were unable to reach or find them. In accordance with how the eligibility status of the rest of the observed addresses in the sample was distributed, one of the 30 addresses was randomly chosen to be ineligible and the remaining 29 to be eligible.

The eligibility rate in the HFCS sample ultimately came to 95% and the nonresponse rate of the eligible households amounted to 50.2%. This means that successful interviews were conducted with 49.8% of the eligible households in the HFCS sample. Just 44.1% of the eligible households actively refused to take part in the survey.

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 chart 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 apartments or in municipalities with higher personal incomes or with higher unemployment rates have a lower probability of participating. On the other hand, households contacted by older interviewers exhibit a significantly higher probability of responding than households contacted by younger interviewers. Moreover, households that were contacted by interviewers with a university degree, or households that live in neighborhood without graffiti or in municipalities with a higher average population age had an increased probability of participating. This suggests that nonresponse is not random.

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	0.00446*** (0.000545)
Building characteristics (reference group: detached single-family house)	
Semi-detached single-family house	0.108 (0.166)
Single-family townhouse	-0.0911 (0.190)
Residential farm building	-0.279 (0.178)
Apartment in a (high-rise) apartment building	-0.419*** (0.0797)
Student dormitory/rented room	-0.601 (0.429)
Other type of building	1.827*** (0.444)
Building design characteristics (reference group: premium)	
Very good	-0.0765 (0.187)
Medium	-0.102 (0.189)
Basic	-0.168 (0.205)
Very basic	0.120 (0.287)
Location characteristics (reference group: city center)	
Between the city center and suburbs	0.225*** (0.0796)
Suburbs and city outskirts	0.130 (0.0844)
Countryside	-0.121 (0.104)
Graffiti in the neighborhood (reference group: many)	
Location – graffiti = 2, some	0.879** (0.375)
Location – graffiti = 3, few	0.485 (0.361)
Location – graffiti = 4, none at all	0.601* (0.359)
Sample design variables	
Design weight	0.000915*** (0.000202)
Interviewer characteristics	
Female interviewer	-0.311*** (0.0590)
Interviewer's age	0.0141*** (0.00273)
Interviewer's second-stage tertiary education	0.302*** (0.0677)
Interviewer's working experience in months	-0.00168*** (0.000367)
Variables at the municipality level	
Average per capita income per municipality in 2011	-0.0000299*** (0.0000119)
Share of employees in the primary sector per municipality in 2011	-4.090*** (1.134)
Share of university-trained population per municipality in 2012	-1.199** (0.474)
Unemployment rate per municipality in 2011	-6.096*** (1.408)
Average age of population per municipality in 2011	0.00454*** (0.00221)
Variables at the district level	
Average crime rate per district in 2009 and 2010	0.000160 (0.000360)
Constant	-0.325 (0.609)
Observations ¹	6,023

Source: HFCS Austria 2014, OeNB.

¹ The remaining 285 observations in the dataset are ineligible and are therefore not included in the regression.

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

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 in the sample)

Step 2: These households' response probabilities were grouped into seven 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 pre-specified 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 adjusted R^2 of this OLS regression exceeds 95%. This is the case for seven classes in the second wave of the HFCS in Austria. Finally, the average response propensity for each class was calculated (unweighted total number of respondent households/unweighted total number of households).⁵

⁵ The average response propensity is unweighted (with respect to the design weights) for efficiency reasons. See Little and Vartivarian (2003) for more details.

Table 14

HFCS nonresponse weights by response propensity

Classes	Predicted response propensity %	Weight
I	0 to 33	3.417
II	33 to 41	2.694
III	41 to 48	2.375
IV	48 to 55	1.853
V	55 to 64	1.664
VI	64 to 75	1.517
VII	75 to 100	1.216

Source: HFCS Austria 2014, OeNB.

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.⁶ 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.

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 underrepresented. 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

⁶ 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.

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.⁷

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, during the second wave of the HFCS in Austria, the question the residence is registered as a main residence or not was recorded in the interview, 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 frame population size only for households registered at their main residence according to the microcensus. Households not registered at their main residence are then added.⁸ This increases comparability between the HFCS and the microcensus in 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 computed following the poststratification cell adjustment method (Biemer and Christ, 2008) and using the Austrian microcensus data (2014 Q4) 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 depending on the registration status. For households registered at their main residence, the province, the tenure status of the main residence and household size serve as poststratification variables. No poststratification is performed for all other households, because, as described above, they are not included in the external dataset.

⁷ For the microcensus, see Haslinger und Kytir (2006), p. 512 f; for the EU SILC, see Statistics Austria (2015), p. 45.

⁸ Before the poststratification adjustment, the HFCS frame population encompassed 3,875,337 households, consisting of households registered at their main residence (3,802,620) and households not registered at their main residence (72,717 or about 2% of households). After the poststratification adjustment of households registered at their main residence, the population of these households comes to 3,789,808, which corresponds to the household population according to the 2014 Q4 microcensus. As a result, the final HFCS household population amounts to 3,862,525 (= 3,789,808 + 72,717).

Step 2: Calculate the average propensity to be included in the frame population for each cell:

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

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 propensity for each cell.

Households registered at their main residence were grouped according to size, one group containing households with one to four individuals and one containing those with five or more individuals.⁹ This ensures that larger households are not underrepresented in the HFCS sample. Additionally, households registered at their main residence were grouped into (part) owners¹⁰ and tenants. Moreover, the households were assigned to the nine provinces.

Table 15 shows the HFCS poststratification weights – 36 values, i.e. one value per combination of registration status, province, tenure status of the main residence and household size. The table shows, for example, that large households are underrepresented in the HFCS sampling frame, as they tend to exhibit larger poststratification weights.

Table 15

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

	Official main residence				Other households			
	Household size (number of persons)				Household size (number of persons)			
	1 to 4	5 or more	1 to 4	5 or more	1 to 4	5 or more	1 to 4	5 or more
	Homeowners		Tenants		Homeowners		Tenants	
Vienna	1.068	0.978	0.964	1.394				
Lower Austria	1.418	2.463	0.556	1.945				
Burgenland	0.870	2.425	1.101	0.335				
Styria	0.968	1.251	0.930					
Carinthia	1.115	1.347	0.685	1.608		1		
Upper Austria	1.270	1.367	0.941	1.052				
Salzburg	0.743	0.478	2.330	3.142				
Tyrol	0.956	1.455	1.118	1.635				
Vorarlberg	1.480	0.824	0.408	0.562				

Source: HFCS Austria 2014, OeNB.

7.2.5 Final weights

Three different weights were computed 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 :

⁹ Given the very low number in a poststratification cell, for main residence tenants in Styria the household size cells were aggregated.

¹⁰ Includes free users of main residences.

$$w_i = \frac{W_{Di} \cdot W_{NRi} \cdot W_{PSi}}{I}$$

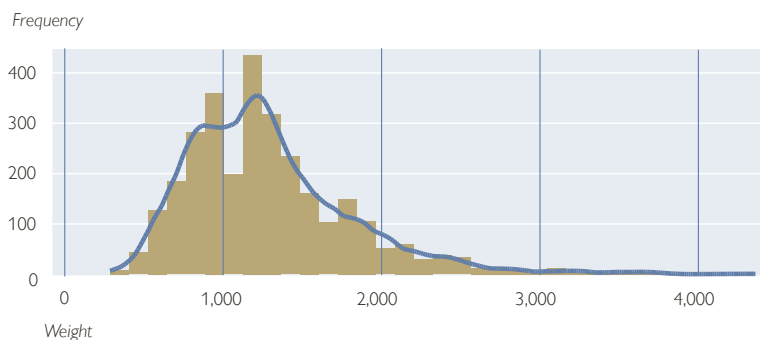
$$w_i = \frac{1}{\text{Prob}(i \text{ is selected}) \cdot \text{Prob}(i \text{ responds} | i \text{ is selected}) \cdot \text{Prob}(i \text{ is included in the frame} | i \text{ is selected and responds})}$$

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

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

Chart 6

Distribution of HFCS final weights



Source: HFCS Austria 2014, OeNB.

The combination of nonresponse weights and poststratification weights results in 252 different weight adjustment cells based on registration status, province, tenure status of the main residence, household size 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 6. The HFCS final weights range from 287 to 4,360, with the mean being 1,289 and the median 1,207.

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. In addition, a slight bipolarity can be identified in the distribution of the weights as a result of the nonresponse weights.¹¹

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 strongly downweighted from 24.7% to 23.1%. This means that despite their high nonresponse rate, overall households in Vienna were clearly overrepresented in the sample with respect to the target population. The comparison also shows that households with higher income or higher net wealth were underrepresented in the unweighted sample, which is probably caused by these households' higher nonresponse rate.

¹¹ Table 14 shows a fairly strong increase in nonresponse weights from class IV to class V, which causes the above-mentioned bipolarity in the distribution of the final weights.

Using 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.

7.4 Concluding remarks

We computed 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.

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 16.7% ($UWE = 1 + \text{coefficient of variation}^2 = 1.167$) as a result of weighting. The adapted sample design made it possible to further improve this value compared to the first wave. Therefore, it is 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[®] is provided in chapter 9 (User guide).

Table 16

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

	Mean	
	Unweighted	Weighted
Household size (number of persons)	2.07	2.14
% of households		
Vienna	24.7	23.1
Lower Austria	16.6	18.9
Burgenland	3.2	3.2
Styria	13.8	13.7
Carinthia	6.3	6.5
Upper Austria	15.9	15.9
Salzburg	7.2	6.1
Tyrol	7.8	8.4
Vorarlberg	4.3	4.2
EUR		
Estimated household monthly net income	2,388	2,450
Household net wealth	227,887	258,414

Source: HFCS Austria 2014, OeNB.

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, calculating 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 overrepresent 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, literature preferred linearization to replication, as it requires less computational power. However, an important disadvantage of linearization is that data protection regulations prevent the required information necessary for linearization from being provided. One way to avoid the problem that certain information is not available for privacy reasons is to use replicate weights. Since replicate weights consist of many variables and their values are based on information not provided to the user of the dataset – e.g. stratum and primary sampling unit (PSU) variables – it is not possible for the data user to identify a given respondent (see Stata Library, 2016).

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., 2010).

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 decided that the variance estimation method to be employed for the HFCS should be based on the use of replicate weights.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. (2010).

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 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 $r=1, \dots, R$ are 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 cases, the estimators of the variance of all replication methods converge toward 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., 2010) 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 in 50 out of 185 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 downward biases of standard errors, which lead to results that are too often considered statistically significant.

Table 17 shows how stratum size (in terms of the 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 185 to 135, which means stratification is still very high. Moreover, the mean stratum size increases from 3.3 PSUs to 4.6 PSUs 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 in 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 correction³ is performed.

8.2.3 Calibration 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 post-stratification 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 $r = 1, \dots, R$ 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 in order 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

Table 17

Comparison of HFCS design strata and HFCS pseudo strata

	Design strata	Pseudo strata
Number of strata	185.0	135.0
Mean size	3.3	4.6
Median size	2.0	2.0
Minimum size	1.0	2.0
Maximum size	37.0	37.0

Source: HFCS Austria 2014, OeNB.

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

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 SSUs are sampled without replacement in the original HFCS sample design (see footnote 3).⁴

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	1,289	1,207	287	4,360	3,862,526
1 st set of replicate weights	1,289	1,040	7	14,374	3,862,526
2 nd set of replicate weights	1,289	989	10	11,418	3,862,526
3 rd set of replicate weights	1,289	1,023	8	10,852	3,862,526
998 th set of replicate weights	1,289	1,104	10	8,369	3,862,526
999 th set of replicate weights	1,289	985	6	11,201	3,862,526
1,000 th set of replicate weights	1,289	974	7	10,349	3,862,526

Source: HFCS Austria 2014, 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

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.

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. 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[®].

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 chapter¹ provides Stata^{®2} code that users can employ step by step to account for all of these features.³ Some extracts of the code were provided by Sébastien Pérez-Duarte⁴ (ECB) and have been slightly altered and expanded for release here. The ECB is expected to also make several program codes available in fall 2016 in addition to publishing the dataset. In this chapter Stata[®] program code is contained in the blue boxes. It can be copied into the Stata[®] command window,⁵ but must be run in the sequence outlined below (altering the sequence and/or including other lines may make the code corrupt). Additionally, the online appendix contains a do-file “user_guide.do” with all the steps that are laid out below.⁶ 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.

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 variables⁷ (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.”

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 13.1 and are not valid for previous Stata[®] versions.

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 Principal Economist Statistician in the Statistics Development/Coordination Division of the ECB.

⁵ Due to the way Stata[®] handles line breaks, they may need to be deleted if the program code is copied by hand.

⁶ 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.

⁷ The ECB is expected to make the definitions of the aggregated variables and the datasets available in fall 2016.

```

*****
***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"

*Merging the p and h files together (wide format)
forvalues i=1(1)5 {
  use "$hfcsdata\P`i'.dta", clear
  drop id hid survey
  foreach var of varlist sa0010- fra0500 {
    local `var'lab: variable label `var'
  }
  reshape wide ra0?0* fra0?0* ra0020 fra0020 ra0030 fra0030 ra0040 fra0040 p* fp* , ///
  i(sa0010 sa0100) j( ra0010)
  foreach j of varlist ra* fra* p* fp* {
    local last2car=substr("`j'", `=length("`j'")-1', 1)
    local last1car=substr("`j'", length("`j'"), 1)
    if "`last2car'"=="1" {
      local firstcar=substr("`j'",1, `=length("`j'")-2')
      rename `j' `firstcar'_'last2car'_'last1car'
      label variable `firstcar'_'last2car'_'last1car' ///
      "`firstcar'lab' - `last2car'_'last1car'"
    }
    else {
      local firstcar=substr("`j'",1, `=length("`j'")-1')
      rename `j' `firstcar'_'last1car'
      label variable `firstcar'_'last1car' "`firstcar'lab'-`last1car'"
    }
  }
  save "$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 "$hfcsdata\templ.dta", clear
forvalues j=2(1)5 {
  append using "$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 multiply 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 naming of the P-file variables, and by merging the resulting dataset with the H-files. They are provided in wide format⁸ (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.

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 estimate` 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.⁹ 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 as “unregistered varying.”

⁸ 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 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.”

```

*****
***Preparing the data for mi import
*****

*Create the zero implicate to simulate the original data
*Use one implicate of the data
use "$hfcsdata\templ.dta", clear
*Replace the implicate number by "0" to simulate the original data
replace im0100=0
*Append all other implicates
append using "$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* 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=""
foreach var of varlist hb* hc* hd* hg* hh* hi* 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* {
  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.¹⁰ The `esampvaryok` option can be useful when the sample size varies across implicates due to imputations.¹¹ Stata[®] versions below Stata[®] 12 do not allow the use of replicate weights together with multiply imputed data. For Stata[®] 12 or higher, 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 (i.e. `"...svy, subpop(dummy)..."`).¹² 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
```

¹⁰ 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.

¹¹ Rubin's combination rules (see e.g. Little and Rubin, 2002) 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.

¹² The use of an `if`-condition does not account for the uncertainty of the subsample size and therefore yields incorrect variance estimators.


```

*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)

```

9.6 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 since 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)

*10th percentile 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. The Stata[®] code in the online appendix will be updated as required, to include program codes for other HFCS-relevant topics. Every additional do-file will be supported with the corresponding documentation.

10 Changes from the first to the second wave of the HFCS

10.1 Introduction

The HFCS has now been conducted twice in Austria. The field phase of the first wave lasted from September 2010 to May 2011. In the second wave, the fieldwork was carried out between June 2014 and February 2015. The second wave drew heavily on the experience gained in the first wave. After all, HFCS data are used for a broad range of research and by all relevant institutions in Austria.

This chapter offers a short, but comprehensive insight into the changes made between the two HFCS waves for those who already have experience with evaluating data from the first wave in Austria. Furthermore, this chapter provides the foundation for evaluations based on both waves of the HFCS in Austria, which require an understanding of how the two survey waves differ.

The structure of this chapter mirrors the structure of the documentation as a whole: Following an overview of the key changes made to the questionnaire (section 10.2) and to interviewer training and selection (section 10.3), we discuss editing measures (section 10.4), the multiple imputation process (section 10.5) and the sampling design (section 10.6). The final two sections deal with the construction of survey weights (section 10.7) and replicate weights (section 10.8). The user guide (see chapter 9) is not discussed here, since it was left broadly unchanged. The chapter is completed with the concluding remarks.

10.2 Questionnaire

10.2.1 Recording of household matrix

The second wave benefited from a more efficient, and hence faster, technical process for recording the demographics for all household members and identifying the financially knowledgeable person (and if necessary reference person).

Once interviewers had established the number of household members, a matrix appeared which allowed them to record the basic data for all people in the household: name (required for referencing while completing the questionnaire), age, gender, relationship with the reference person and financial affiliation with the household. This information was subsequently shown in the form of a list, to enable respondents to cross-check the data, with a view to making adjustments, deleting a person or adding another person. The interview did not start until the basic data for all household members had been confirmed. Compared with the first wave, this means that above all the questions about age and gender were moved from the first chapter of the survey to the household matrix, which helped save time.

10.2.2 Lists of predefined ranges

As outlined in section 2.6.2, all questions about amounts of money were asked in loops. If respondents were unable or unwilling to indicate specific amounts or individual ranges, they were asked to choose a range from a list of predefined ranges. While the same list of predefined ranges was used for all questions in the first wave, three different lists were used in the second wave, depending on the question at hand (see table 19).¹ The first column shows the list of ranges used in the first wave, and the other three columns the lists of ranges (A, B and C) used in the second wave.

¹ See the online appendix for the showcards used during the interviews.

Table 19

Different lists of predefined ranges in the first wave and the second wave (in EUR)

First wave (HFCS Austria 2010)	Second wave (HFCS Austria 2014)		
List of predefined ranges	List A	List B	List C
<i>EUR</i>			
A 1 – below 101	A 1 – below 101	A 1 – below 10,001	A 1 – below 1,001
B 101 – below 501	B 101 – below 201	B 10,001 – below 50,001	B 1,001 – below 2,501
C 501 – below 1,001	C 201 – below 301	C 50,001 – below 75,001	C 2,501 – below 5,001
D 1,001 – below 2,501	D 301 – below 401	D 75,001 – below 100,001	D 5,001 – below 7,501
E 2,501 – below 5,001	E 401 – below 501	E 100,001 – below 150,001	E 7,501 – below 10,001
F 5,001 – below 7,501	F 501 – below 751	F 150,001 – below 200,001	F 10,001 – below 15,001
G 7,501 – below 10,001	G 751 – below 1,001	G 200,001 – below 300,001	G 15,001 – below 20,001
H 10,001 – below 25,001	H 1,001 – below 1,501	H 300,001 – below 400,001	H 20,001 – below 25,001
I 25,001 – below 50,001	I 1,501 – below 2,001	I 400,001 – below 500,001	I 25,001 – below 30,001
J 50,001 – below 75,001	J 2,001 – below 3,001	J 500,001 – below 750,001	J 30,001 – below 35,001
K 75,001 – below 100,001	K 3,001 – below 5,001	K 750,001 – 1 million	K 35,001 – below 40,001
L 100,001 – below 250,001	L 5,001 – below 7,501	L more than 1 million – 3 million	L 40,001 – below 50,001
M 250,001 – below 500,001	M 7,501 – below 10,001	M more than 3 million – 5 million	M 50,001 – below 75,001
N 500,001 – 1 million	N 10,001 – below 25,001	N more than 5 million – 10 million	N 75,001 – below 100,001
O more than 1 million – 5 million	O 25,001 – below 50,001	O more than 10 million	O 100,001 – below 200,001
P more than 5 million – 10 million	P more than 50,000		P 200,001 – below 300,001
Q more than 10 million – 25 million			Q 300,001 – below 500,001
R more than 25 million – 50 million			R 500,001 – 1 million
S more than 50 million – 100 million			S more than 1 million
T more than 100 million			

Source: HFCS Austria 2010 und 2014, OeNB.

Table 20

Unweighted percentiles of selected variables in the HFCS Austria 2010 wave

Percentiles	Food-at-home consumption ¹	Current value of main residence ²	Gross income from dependent employment ³
<i>EUR</i>			
P10	170	80,000	6,400
P20	200	113,000	11,100
P30	250	148,000	14,400
P40	300	169,000	17,500
P50	350	200,000	20,200
P60	400	231,000	24,000
P70	450	275,000	28,600
P80	500	342,000	34,800
P90	600	485,000	45,500

Source: HFCS Austria 2010, OeNB.

¹ Rounded to the nearest EUR 10.² Rounded to the nearest EUR 1,000.³ Rounded to the nearest EUR 100.

The new lists were created based on the (unweighted) empirical distribution of the euro measures recorded for the corresponding variables.² See table 20³ for the distribution of amount spent on food at home (as an example of list A), the distribution of the value of main residences at the time of the interview (as an example of list B) and the distribution of gross income from dependent employment (as an example of list C).

With some exceptions, all empirically observed distributions were found to fall into very few (mostly one or two) ranges from the list of ranges used in the first wave. The switch to three different lists of ranges in the second wave allowed us to capture a much more accurate range in which a given amount

² The average of all implicates was used.³ This table was first published in Lindner et al. (2014).

falls. For instance, while in the first wave the values for the household main residence at the time of the interview were essentially spread over three ranges from the list of ranges, they are now spread over six ranges. Conversely, those ranges that were less relevant for a particular question, for instance values of below EUR 10,000 for real estate property, were merged into one range. Respondents were also asked to confirm all predefined ranges that they had selected, as well as all amounts or individual ranges that they had specified (see section 2.6.2).

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 income. Questions about financial assets were either allocated to list A or list C, depending on the distribution of assets as observed in the first wave of the survey.⁴ The predefined ranges referred to amounts in euro only.

10.2.3 Recording households owning a farm

The first wave of the HFCS showed that farmers found it particularly challenging to break down their assets so that they would fit the structure defined by the HFCS. In the first HFCS wave, support in this respect was limited to a few notes in the questionnaire and the training of interviewers. For the second wave, the questionnaire was improved for the group of farmers as outlined below (and in section 2.6.3).

Before the interview started, respondents were classified by the interviewer as running an “Agricultural business” or as running “No agricultural business.” The classification was straightforward in all but a few cases. Even if interviewers misclassified a respondent, they still recorded all the relevant information.

Specifically, extra information was recorded on farmers as follows:

- Was it possible to differentiate housing assets (i.e. the household main residence) from business assets? [in the questionnaire chapter on the household main residence]
- If not, what percentage of the recorded value did respondents allocate to their main residence? [in the questionnaire chapter on the household main residence]
- Does the value recorded for investment in a self-employment business include the recorded main residence? [in the questionnaire chapter on investment in a self-employment business]

Moreover, for questions relating to the value of the main residence, the yes/no question on properties other than household main residence, as well as the question about investment in a self-employment business and its value, farmers received detailed instructions as to which components of their household balance sheet were to be recorded under which position.

This approach greatly facilitated in particular subsequent imputations (see also section 5.4.5) with a view to breaking down the assets into farmers’ main residences and their agricultural businesses. In addition, the training of interviewers was improved to enable them to better handle interviews with farmers (see section 10.3).

⁴ See the questionnaire in the online appendix for a detailed overview of which ranges (list A, B or C) were used for which questions.

10.2.4 Loans from relatives and friends

For the second wave, the recording of loans received from relatives and friends was adjusted in line with international standards. In the first wave, information on such liabilities had been collected within the set of questions about other uncollateralized loans. Information on private loans was collected from the question regarding which institution the loan was taken out with, which respondents could answer with “family” and “friends.”

The first wave of the HFCS showed that loans among private individuals are a significant component of the household balance sheet and that less information is required to record private loans from other lenders. Therefore, the decision was made to record these two categories separately. As a result, the second wave of the HFCS in Austria utilized two loops⁵ with up to three iterations, for both liabilities to “relatives and friends” and liabilities to other institutions (“other uncollateralized loans”).

10.2.5 Recording inheritances and gifts

The results of the first HFCS wave (see e.g. Fessler and Schürz, 2015) show that transferring capital in the form of inheritances and gifts is a major channel through which households accumulate wealth. The HFCS addresses capital transfers in two sections. First, respondents are asked to indicate how they acquired the (part-) ownership of their main residence. Second, information about ownership transfer of all other wealth is collected using loops. In the second HFCS wave in Austria this loop is repeated up to five times (compared to three in the first wave). This was also necessary since in the loop the value of a gift/inheritance at the time of ownership transfer is collected, which makes collection of this information in summary questions difficult. However, the core dataset published by the ECB so far contains only three iterations of the loop question on inheritances and gifts.

10.2.6 Comment fields in the questionnaire

In the case of questions which proved difficult for respondents or were identified as essential in the first wave, the digital questionnaire for the second wave was expanded to include an additional box, allowing both numeric values and text to be entered. The information recorded in those boxes proved to be very relevant during the postinterview checks. Such explanations have often helped solve problems recording some information that would otherwise have required follow-up phone calls (see section 4.4.2).

10.3 Interviewers

With regard to the interviewers, the key improvements observed can be attributed to the experience interviewers gained in the first wave. About half of the interviewers employed had already worked on the HFCS in 2010/11. Moreover, the training for interviewers in the second wave was revised to accommodate the experience gained from the first wave. In particular, rather than offering separate theoretical and practical training sessions, the two components were integrated more closely. After an initial theory session, interviewers had a chance to run through some items of the questionnaire in a practical, interactive session. There-

⁵ See section 2.6.1. on the structure and navigation of loops.

after, the training alternated between theory and practice. A mock interview session using a highly complex household owning an agricultural business as an example made sure that all aspects of the questionnaire were practiced and discussed in detail.

In the first wave interviewers who joined the interview team during fieldwork were trained in additional sessions conducted by the survey company on its own, without involving an OeNB HFCS staff member. The quality of interviews conducted by these interviewers was poorer than those of interviewers who had attended the regular training. Therefore additional training sessions by the survey company were ruled out for the second wave.

10.4 Consistency checks and editing

The consistency checks were intensified further in the second wave. On the one hand, the experience gained allowed for greater efficiency in checking, and on the other hand it was possible to increase the number of checks. In the second wave, the number of consistency checks programmed into the digital questionnaire was increased significantly.

The four-eyes principle for case-by-case reviews was retained; the answers of all households (whether interviewed successfully or not) were checked for internal consistency, grouped by interviewers. Moreover, the higher frequency with which the survey company forwarded data to the OeNB during the field phase further diminished the gap between the interviews and potential follow-up queries.

The inclusion of verbatim record fields for complex questions, as well as the expansion of the comment fields in which interviewers made postinterview comments on all households, facilitated the ex post editing measures. The verbatim records collected helped clarify numerous problems (see also chapter 4). Nonetheless, if required ex post queries were made by phone.

10.5 Multiple imputations

A major change in the HFCS imputation procedure made for the second wave was not to conduct any weighted regressions. While in the first wave weighted regressions were estimated using the final survey weights (in step 3 of the procedure), those weights were only used as model predictors in the second wave. This approach is in line with the current trend in imputation literature (see e.g. Frumento et al., 2012): The purpose of multiple imputations is to produce good forecasts of the missing values (and the corresponding degree of uncertainty); the weighting of households should not occur before the final analysis of the dataset and a general assessment of the population have been carried out. This change made for the second wave ought to help decrease the standard errors slightly, because the variance of the imputed values within every multiple imputation sample, as well as the variance of the imputed samples across the multiple imputation sample, is relatively small as a result of the nonweighting of regressions.

Further improvements of the HFCS imputation procedure are aimed at enhancing the consistency across the variables and improving the convergence as well as the evaluation of the convergence of the imputation model. We were able to increase the consistency between some quantitative variables by imputing the variables as a share of other variables (see section 5.4.5) rather than imputing each variable separately. As a case in point, outstanding consumer loans were imputed

as a share of the original size of the loan; or the size of a household's first consumer loan was imputed as a share of all loans taken out by the specific household.

The convergence of the HFCS imputation procedure was potentially improved in the second wave by increasing, from 6 to 10, the number of cycles which define how often regressions are to be re-estimated and how often imputed values are to be updated. The higher this figure is, the closer the imputed values should be to a draw of the common predictive a posteriori distribution of the variables with missing values (see section 5.4.9). As discussed in chapter 5, this relationship has been cross-checked in general in simulations, but it is yet to be underpinned with theoretical evidence (see section 5.3).

The criteria for evaluating convergence were expanded in the second wave to include not only graphical checks, but also the widely used Gelman-Rubin criterion (see e.g. Cowles and Carlin, 1996). Care is taken to keep the variance of the mean across the multiple imputation sample relatively small compared with the variance of the mean across the cycles (see section 5.4.9).

10.6 Sampling

Essentially, the sampling design was improved in three ways for the second wave of the HFCS in Austria: The gross samples were enlarged, the clusters of municipalities (other than Vienna) with more than 50,000 inhabitants were adjusted in size, and changes were made to the selection probability for PSUs.

In particular, the gross sample in Austria was increased from 4,436 to 6,308 households. As a result, it was possible to arrive at a significantly higher number of successfully interviewed households. Therefore the dataset generated by the second wave of the HFCS in Austria contains 2,997 households that were interviewed successfully (compared with 2,380 in the first wave). A larger sample should increase the precision of the estimators based on the HFCS and further facilitate the analysis of subgroups, given the larger amount of interviews that could be analyzed.

In the first wave, households were clustered to groups of 8 per PSU in Vienna (as opposed to 12 in the rest of Austria) to account for the smaller geographical distances between households in Vienna. In the second wave, this strategy was applied to all strata with municipalities of 50,000 and more inhabitants. Thus, all other major cities in Austria⁶ now have a sampling design with a relatively small number of SSUs (i.e. households) per drawn PSU (i.e. enumeration district). The theoretical basis of the sampling suggests that this strategy reduces the variance of individual estimators⁷ generated on the basis of the survey.

The biggest and most important change compared with the first wave is the improvement with regard to the selection probability for PSUs. In the first wave, each PSU within a stratum had the same probability of being drawn. In the second wave, the selection probability was tied to the number of households in a given PSU and was defined by the ratio of the number of households in a given PSU divided by the number of households in this stratum times the number of drawn PSUs. In the second wave the PSUs were also drawn with replacement. Some

⁶ Vienna, St. Pölten, Graz, Klagenfurt, Villach, Linz, Wels, Salzburg and Innsbruck.

⁷ See also chapter 8 in Levy and Lemeshow (2008).

PSUs could therefore be drawn more than once. This adjustment was made with a view to decreasing the variance of household weights (see also section 10.7).

The sampling design also changed slightly as a result of changes in the sampling frame resulting from mergers of municipalities or enumeration districts. The sampling was based on external data provided by Statistics Austria and the Austrian Post Office at the time of sampling. Following the merger of some municipalities in the time between the first and the second wave, enumeration districts underwent changes, as did the strata allocation of some municipalities. However, these changes should play a minor role in the evaluation of the data.

10.7 Construction of survey weights

The adjustments made to the sampling design also affected the calculation of the design weights. As in the first wave of the HFCS in Austria, the design weights were also based on the selection probability of households in the second wave, with the design weights being the inverse of the selection probabilities. For a detailed overview of the calculation of design weights, see section 7.2.1. The adjustments made to improve sampling and the resultant design weights influenced the distribution of the weights in the second wave.

See chart 7 and table 21⁸ for a comparison of the distribution of design weights in the first and the second waves of the HFCS in Austria.

Table 21

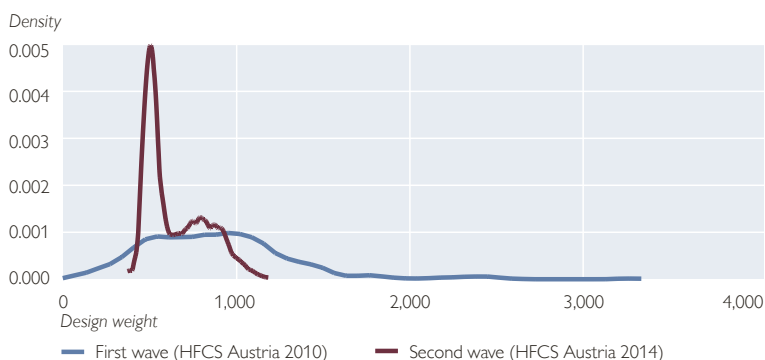
Comparison of descriptive statistics of design weights in HFCS Austria 2010 and 2014

	First wave (HFCS Austria 2010)	Second wave (HFCS Austria 2014)
Minimum	61	369
Median	857	553
Mean	884	642
Maximum	3,271	1,183
Standard deviation	434	177
Number of observations	4,436	6,308

Source: HFCS Austria 2010 und 2014, OeNB.

Chart 7

Comparison of the distribution of design weights in HFCS Austria 2010 and 2014



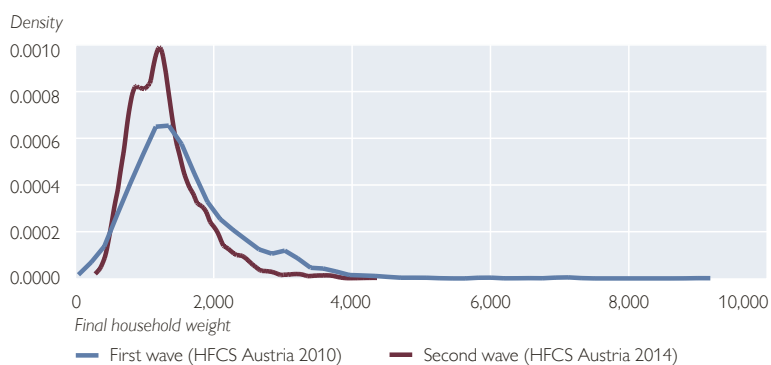
Source: HFCS Austria 2010 und 2014, OeNB.

It can be seen that the variance of the design weights has diminished significantly in the second wave due to the adjustments made to the sampling design. In particular, the standard deviation dropped from about 430 in the first wave to about 180 in the second wave. Furthermore, the range of the design weights was reduced significantly, to values between 369 and 1,183, in the second wave. As a result of the higher number of households in the gross sample, the mean of the design weights decreased from about 880 to about 640.

⁸ Both chart 7 and table 21 were published and discussed in Lindner et al. (2014).

Chart 8

Comparison of the distribution of final household weights in HFCS Austria 2010 and 2014



Source: HFCS Austria 2010 und 2014, OeNB.

Table 22

Comparison of descriptive statistics of final household weights in HFCS Austria 2010 and 2014

	First wave (HFCS Austria 2010)	Second wave (HFCS Austria 2014)
Minimum	169	287
Median	1,429	1,207
Mean	1,586	1,289
Maximum	9,054	4,360
Standard deviation	834	527
Number of observations	2,380	2,997

Source: HFCS Austria 2010 und 2014, OeNB.

However, the adjustments are aimed at enhancing the final household weights, which goes hand in hand with changes in the design weights, where the changes appear to be more prominent. For details on the changes of the final household weights, see chart 8 and table 22.

The enhanced sampling procedure is also evident from the final household weights. In particular, the range of the final household weights was reduced from a range of about 170 to 9,050 in the first wave to about 290 to 4,360 in the second wave. It should be noted that nonresponse and poststratification adjustments cause the variability of the final household weights to be higher than that of the design weights.

The nonresponse adjustment method was also improved for the second wave. In both waves, this step in calculating the weights (see also section 7.2.3) is based on a nonlinear model designed to explain nonresponse. For the second wave, the model was expanded to include additional paradata on the interview, the place of residence and residential area, as well as data on the municipalities and political districts. Moreover, the response propensity classes applied were selected in line with the method proposed by Haziza and Beaumont (2007) so as to achieve an optimal tradeoff between bias and variance. Unlike the first wave, where households' response propensities were designed to be grouped into five classes (quintiles), the method used in the second wave is based on seven classes.

In the poststratification process (see section 7.2.4) additional information regarding main residences and second homes of the participating households allowed for a more granular poststratification of households officially registered at their main residence⁹ on the basis of the relevant external data. Whereas in the first wave poststratification with the microcensus data from Statistics Austria was based on household size and municipality size, in the second wave it was based on province, household size and the tenure status of the household's main residence. This made it possible to poststratify the data, not only based on the shares of households in the individual groups, but also based on the number of households in those groups.

⁹ Households not officially registered at their main residence in the centralized residence registry were excluded from the poststratification process, i.e. in the absence of relevant external data they were given a poststratification weight of 1.

10.8 Construction of replicate weights

Like in the first wave, the construction of replicate weights mirrors the computation of the final household weights.

The sampling error calculation model mimics the original sampling procedure as closely as possible. As in the first wave, all single-PSU strata are paired with a neighboring stratum. However, in the second wave, we made sure to pair single-PSU strata with the smaller of two or more neighboring strata (in terms of numbers of PSUs) wherever possible, in order to limit the number of aggregations required. As a result only 50 of the 185 strata had to be aggregated (compared with 81 out of 170 strata in the first wave). Aggregation is necessary as the construction of replicate weights requires at least two PSUs per stratum. With-replacement sampling of PSUs in the sampling error calculation model mirrors the sampling design and is also more in line with survey theory. However, for the purpose of the sampling calculation model we considered all PSUs within a stratum to have the same selection probability.

By enhancing the sampling procedure (and the way it is represented in the sampling error calculation model), we were able to produce more efficient replicate weights in terms of variance. Table 23 shows a comparison of the two survey waves.

Table 23

Comparison of selected replicate weights in HFCS Austria 2010 and 2014

	Minimum	Median	Mean	Maximum	Standard deviation	Number of observations
HFCS Austria 2010						
1 st set of replicate weights	5	1,101	1,586	11,805	1,875	2,380
2 nd set of replicate weights	4	1,089	1,586	14,345	1,941	2,380
3 rd set of replicate weights	4	948	1,586	18,429	2,091	2,380
998 th set of replicate weights	3	1,174	1,586	22,191	2,005	2,380
999 th set of replicate weights	5	1,170	1,586	17,956	2,132	2,380
1,000 th set of replicate weights	3	1,122	1,586	14,139	2,028	2,380
HFCS Austria 2014						
1 st set of replicate weights	7	1,040	1,289	14,374	1,519	2,997
2 nd set of replicate weights	10	989	1,289	11,418	1,472	2,997
3 rd set of replicate weights	8	1,023	1,289	10,852	1,436	2,997
998 th set of replicate weights	10	1,104	1,289	8,369	1,385	2,997
999 th set of replicate weights	6	985	1,289	11,201	1,410	2,997
1,000 th set of replicate weights	7	974	1,289	10,349	1,473	2,997

Source: HFCS Austria 2010 und 2014, OeNB.

It is evident that the standard deviation could also be reduced with regard to the replicate weights. At the same time, the range (minimum to maximum) of individual replicate weights can differ between the two waves for randomly drawn PSUs per resample, meaning there is no clear sign of a significant reduction in this respect.

10.9 Concluding remarks

This chapter provided a small but comprehensive overview of the differences between the first and second waves of the HFCS in Austria. For more detailed information on specific aspects of this publication, please see the relevant chapters or sections in this documentation.

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. 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. 93–109.
- Albacete, N. and M. Schürz. 2013a.** Vergleich der Einkommensmessung für Haushalte in Österreich: HFCS versus EU-SILC. In: Statistiken – Daten und Analysen Q2/13. 88–89.
- Albacete, N. and M. Schürz. 2013b.** Interviewereffekte beim HFCS Austria 2010. In: Statistiken – Daten und Analysen Q3/13. 57–68.
- Albacete, N. and M. Schürz. 2014a.** Paradata im HFCS Austria 2010 – Teil 1: Evaluierung von Non-Response-Fehlern. In: Statistiken – Daten und Analysen Q1/14. 81–97.
- Albacete, N. and M. Schürz. 2014b.** Paradata im HFCS Austria 2010 – Teil 2: Evaluierung von Messfehlern. In: Statistiken – Daten und Analysen Q3/14. 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. 55–63.
- 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.
- 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. 95–110.
- 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.
- 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.
- ECB. 2011.** Core output variables catalogue. www.ecb.int/home/pdf/research/hfcn/core_output_variables.pdf?c6a87a29f0c1cdf4b92526aceef3efea (accessed on December 9, 2016).

- ECB. 2013a.** The Eurosystem Household Finance and Consumption Survey: Methodological report for the first wave. Statistics Paper Series No. 2. April.
- ECB. 2013b.** The Eurosystem Household Finance and Consumption Survey: Results of the first wave. Statistics Paper Series No. 2. April.
- Fessler, P. and M. Schürz. 2013.** Cross-country comparability of the Eurosystem Household Finance and Consumption Survey. In: *Monetary Policy & the Economy* Q2/13. 29–50.
- Fessler, P. and Schürz, M. 2015.** Private wealth across European countries: The role of income, inheritance and the welfare state. ECB Working Paper Series 1847.
- 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. 23–62.
- 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. 76–87.
- 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. Doi: 10.1111/j.1751-5823.2006.00002.x.
- Heeringa, S. G., B. T. West and P. A. Berglund. 2010.** Applied survey data analysis. 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 the SCF panel data. Prepared for the Joint Statistical Meeting in Miami, Florida. August 3.
- 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. and M. Schürz. 2015.** Varianten der Messung von Haushaltsvermögen im HFCS in Österreich. In: *Statistiken – Daten und Analysen* Q2/15. 52–70.
- Lindner, P., M. Schürz. and J. C. Zhan. 2014.** Methodische Verbesserungen im HFCS. In: *Statistiken – Daten und Analysen* Q4/14. 71–83.

- Little, R. J. A. and D. B. Rubin. 2002.** *Statistical analysis with missing data*. 2nd 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*. OeNB. December.
- 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. 2015.** *Standard-Dokumentation Metainformationen (Definitionen, Erläuterungen, Methoden, Qualität) zu EU-SILC 2009*.
- Stata Library – Replicate weights. UCLA:** Academic Technology Services. Statistical Consulting Group. www.ats.ucla.edu/stat/stata/library/replicate_weights.htm (accessed on December 9, 2016).
- 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. 49–64.

Publisher and editor	Oesterreichische Nationalbank Otto-Wagner-Platz 3, 1090 Vienna PO Box 61, 1011 Vienna, Austria www.oenb.at oenb.info@oenb.at Phone (+43-1) 40420-5555 Fax (+43-1) 40420-045499
Editorial board	Ernest Gnan, Doris Ritzberger-Grünwald, Helene Schuberth, Martin Summer
Managing editor	Claudia Kwapil
Editing	Dagmar Dichtl, Ingrid Haussteiner, Ingeborg Schuch
Translations	Owen Davies, Dagmar Dichtl, Jennifer Gredler, Ingrid Haussteiner, Rena Mühldorf, Ingeborg Schuch, Susanne Steinacher
Layout and typesetting	Walter Grosser, Andreas Kullerschitz
Design	Information Management and Services Division
Printing and production	Oesterreichische Nationalbank, 1090 Vienna

DVR 0031577

© Oesterreichische Nationalbank, 2016. All rights reserved.

May be reproduced for noncommercial, educational and scientific purposes provided that the source is acknowledged.

Printed in accordance with the Austrian Ecolabel guideline for printed matter.

Please collect used paper for recycling.

EU Ecolabel: AT/028/024

