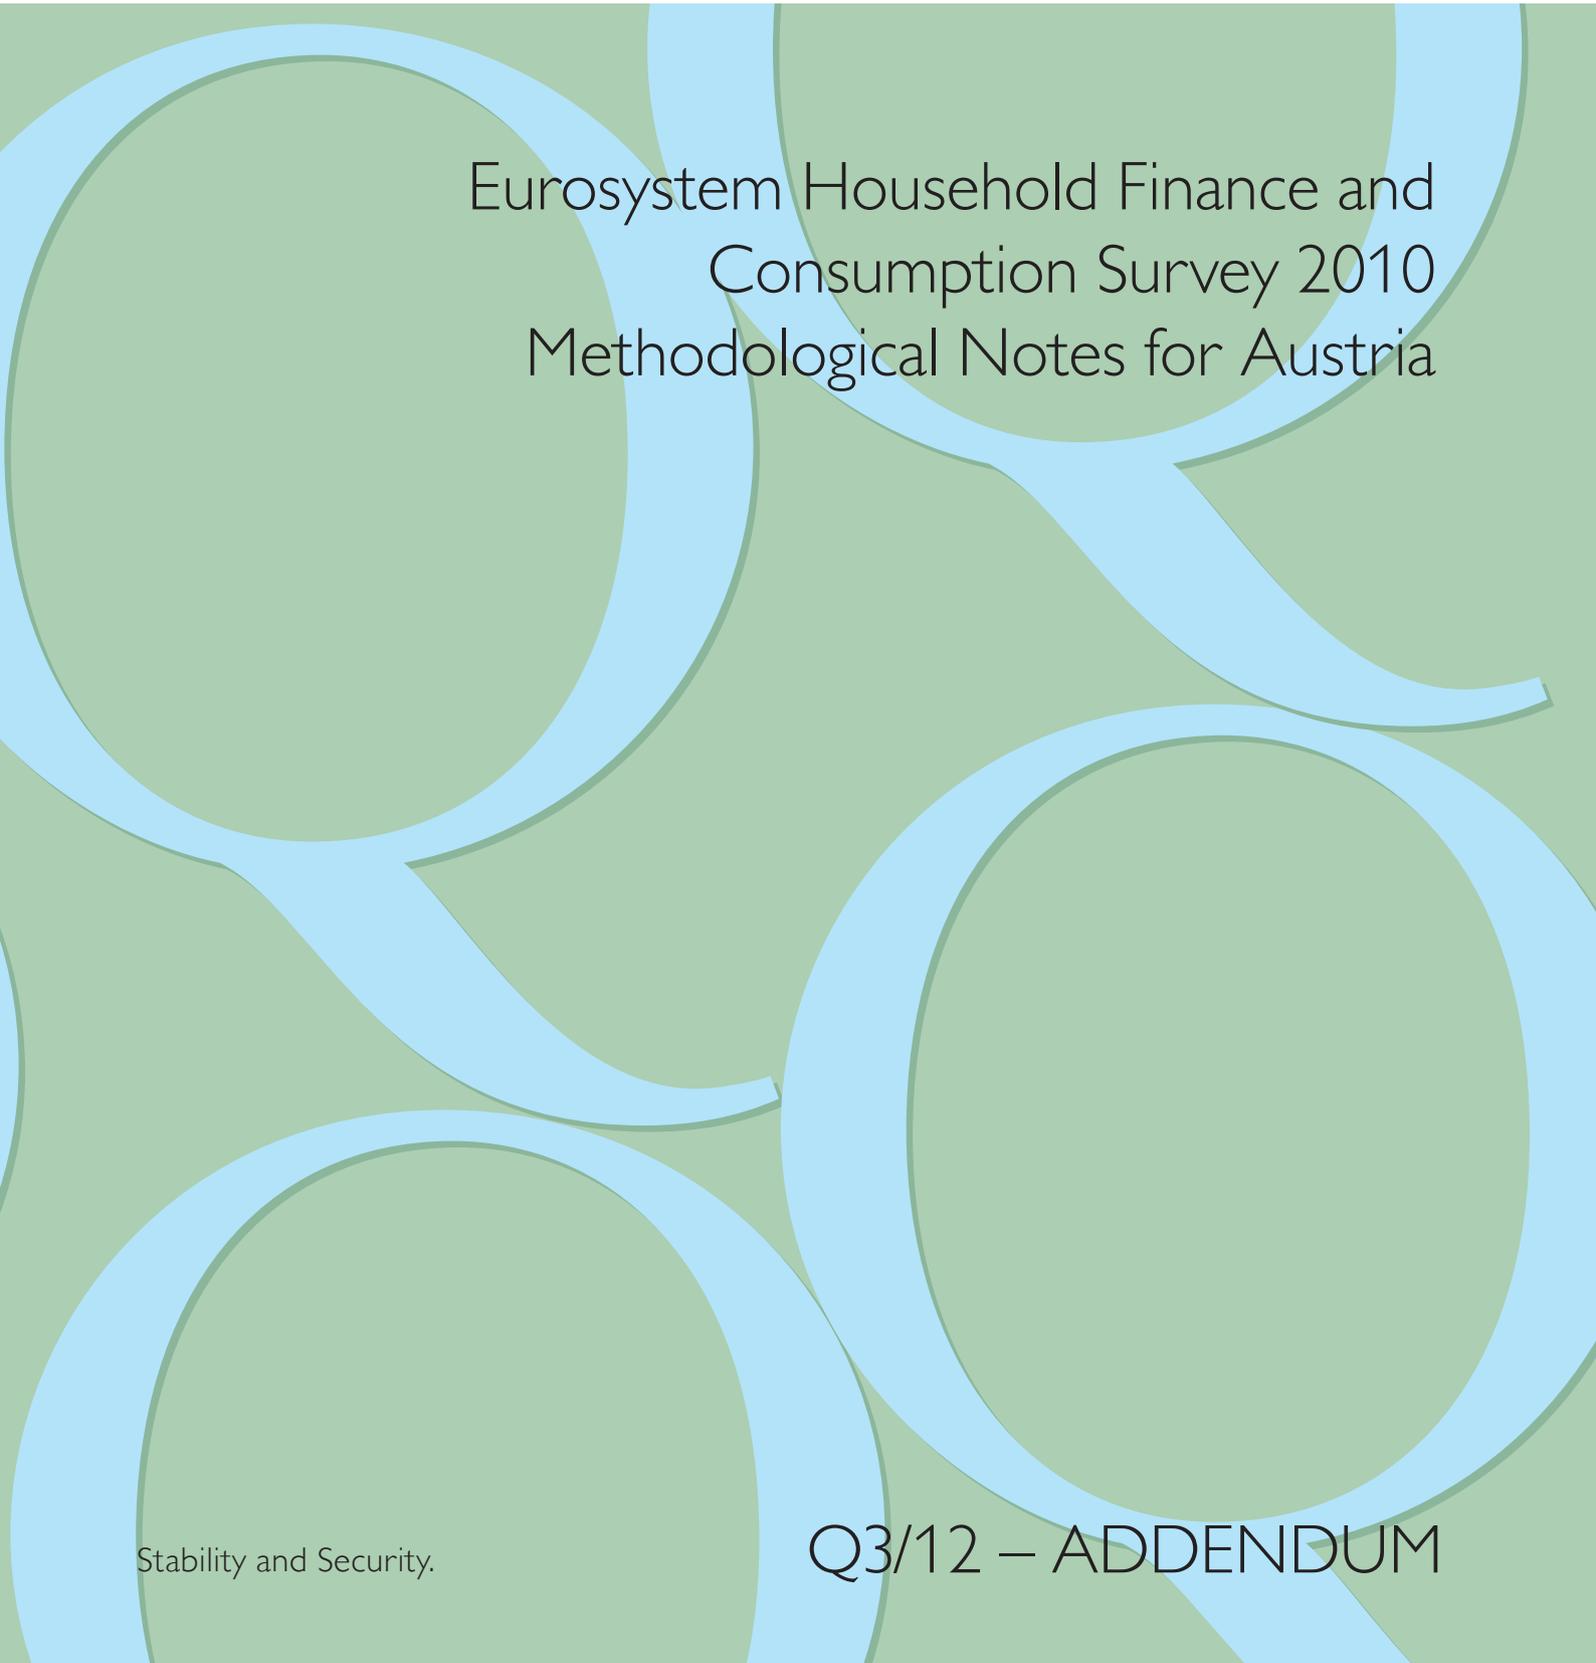


MONETARY POLICY & THE ECONOMY

Quarterly Review of Economic Policy

The background of the lower half of the page is a solid light green color. Overlaid on this are several large, light blue, stylized circular and curved shapes that resemble interconnected loops or a network, creating a modern, abstract design.

Eurosystem Household Finance and Consumption Survey 2010 Methodological Notes for Austria

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	16
2.7 Documents Provided for the Interviewer	19
2.8 Participant Countries	19
2.9 Online Appendix	20
3 Interviewers	21
3.1 The Interviewers' Role in the Survey Process	21
3.2 General Information	21
3.3 Interviewer Training	21
3.4 Contact Strategies and Contact Provisions	23
3.5 Documents and Other Supporting Material	24
3.6 Monitoring	26
3.7 Problems Relating to Interviewers	26
3.8 Online Appendix	27
4 Consistency Checks and Editing	28
4.1 Introduction	28
4.2 Number and Type of Edits	28
4.3 Checks for Consistency during the Interviews	29
4.4 Postinterview Checks for Consistency	30
4.5 Flags	32
4.6 Ex Post Editing	35
4.7 Formatting and Editing after Multiple Imputations	44
4.8 Concluding Remarks and Online Appendix	45
5 Multiple Imputations	47
5.1 Introduction	47
5.2 Item Nonresponse	47
5.3 HFCS Imputation Procedure	50
5.4 Creating the Imputations	51
5.5 Selected Results	58
5.6 Concluding Remarks	60

6	Sampling	61
6.1	Introduction	61
6.2	Target Population and Sample Frame	61
6.3	Background – The (External) Data Sources	63
6.4	Stratification and Sample Size	65
6.5	The Two Stages of the Random Draw	67
6.6	Concluding Remarks	69
7	Construction of Survey Weights	70
7.1	Introduction	70
7.2	Construction of Survey Weights	71
7.3	Selected Results	78
7.4	Concluding Remarks	78
8	Construction of Replicate Weights for Variance Estimation	80
8.1	Introduction	80
8.2	Construction of Replicate Weights	80
8.3	Concluding Remarks	83
9	User Guide	84
9.1	Introduction	84
9.2	Merging the Data Files	84
9.3	Multiple Imputation	86
9.4	Survey Variables	88
9.5	Standard Estimation Procedures	89
9.6	Using Additional Estimation Procedures	90
9.7	Online Appendix	91
	References	92

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 entire balance sheet of households (including all stock and flow positions) together with socioeconomic characteristics. The sampling unit was the household.

Reference Period

The information about all stock positions and socioeconomic characteristics was collected at the date of the survey (survey period: September 2010 to May 2011). Income-related data refer to the 2009 calendar year.

Geographical Scope

Austria

Sampling

Target population

All households in Austria (irrespective of nationality).

Frame population

Postal addresses of all households in Austria.

Sample 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 422 PSUs and 4,436 SSUs in 170 strata.

Survey Company

Institut für Empirische Sozialforschung GmbH – IFES

Field Work

General information

Period:	September 2010 to May 2011
Number of interviewers:	85
Method of data collection:	Computer-Assisted Personal Interview (CAPI)

Interviewer training

Number of interviewer training units:	5
Duration of an interviewer training unit:	1 day

Trial field phase

Number of trial interviews:	50
-----------------------------	----

Contact strategy

All households received a personalized letter from the governor of the Oesterreichische Nationalbank (OeNB) and an information brochure 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, if required. 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 (the actual value during the field phase was around EUR 15).

In addition, each household had the chance to win travel vouchers (one worth EUR 1,000, five vouchers worth EUR 200 each).

Documents used during interviews

Catalogue of show cards, interviewer handbook, glossary

Interviewer monitoring

The survey company monitored the correct conduct of the interviews by checking (by telephone) around every one in six interviews.

The (anonymized) data from the completed household interviews were forwarded to the OeNB in seven batches during the field phase. Each interview and interviewer was monitored closely. The OeNB had the possibility to withdraw interviewers. Three interviewers who were found to have delivered seriously flawed interviews were withdrawn from the survey.

Follow-up investigations 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 400 households).

Editing and Consistency Checks

Number and type of edits

Number of observations:	around 841,000
Thereof edited observations:	around 22,000 (around two-thirds on the basis of verbatim records)
Share of edited observations:	2.6%

Checks for consistency during the interviews

Number of programmed consistency checks during the interviews:	more than 150
----------------------------------------------------------------	---------------

Post-interview checks for consistency

Analysis of the data by experts for each interview, follow-up investigations by telephone to clarify uncertainties, investigation of outliers and checks on the consistency of the information collected, technical review of filters.

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:	6
Median of the variables with missing values per household:	8
Mean of the variables with missing values per household:	17.3

Sample Size and Response Rate

Number of households in the sample (gross sample):	4,436
Number of successfully interviewed households (net sample):	2,380
Number of households that could not be reached (despite five contact attempts):	75
Number of households that refused to participate:	1,693
Number of households that did not participate for other reasons:	120
Number of addresses whose eligibility was unknown:	6
Number of ineligible addresses:	162
Response rate:	around 56%

Weighting

Adjustments for design weights to obtain final weights: nonresponse and poststratification adjustment

- Method of nonresponse adjustment: weighting class adjustment combined with model-based adjustment
- Method of poststratification adjustment: poststratification cell adjustment

Smallest final weight:	169
Median final weights:	1,429
Mean of final weights:	1,586
Highest final weight:	9,054
Sum of final weights (target population):	3,773,956
Unequal weighting effect:	1.28
No trimming and no normalization of weights	

Variance Estimation

Method: rescaling bootstrap procedure	
Number of resamples:	1,000
Number of pseudo-strata:	81
Adjustment of the design weights to obtain the resampling weights: identical to the adjustment of the design weights to obtain the final weights	
Finite population corrections were applied to all resampling weights.	

1 Introduction

The Eurosystem Household Finance and Consumption Survey (HFCS) is the first compilation of data in Austria capturing real assets, financial assets, debt and expenditures of households in one survey, allowing a comprehensive representation of the balance sheet of Austrian households. The HFCS was an ECB initiative and is carried out in a harmonized way in all euro area countries, with the exception of Ireland and Estonia, by these countries' national central banks.

The ex ante harmonization of the survey and the methods used make the HFCS data comparable across all euro area countries. In Austria, the HFCS was conducted by the Oesterreichische Nationalbank (OeNB) in cooperation with the survey company Institut für empirische Sozialforschung GmbH (IFES). The ECB will make the HFCS data of all euro area countries available for research purposes from 2013 on. This publication, entitled *Eurosystem Household Finance and Consumption Survey 2010: Methodological Notes for Austria*, provides an in-depth view of the data collection process and the methods applied. It aims at making the process of data set compilation as transparent as possible and serves as the basis for the correct evaluation using the HFCS data. 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 or material aspects discussed in other chapters. The sequence of chapters is based on the logical relationship between the issues discussed. Closely related topics (e.g. the construction of weights and the correct variance estimation using HFCS data) are arranged in a way to ensure comprehensibility. To avoid redundancies, only the most important information was repeated if deemed necessary. The following eight chapters provide a detailed explanation of each step applied in the survey process.

Chapter 2 on the *Questionnaire* of the HFCS in Austria first explains the content of the survey and then discusses both the individual parts of the questionnaire and the sequence of questions as well as the sampling unit and the particularities of the HFCS questionnaire.

Chapter 3 looks into the role of the *Interviewers* who conducted the interviews in person. Great importance was attached to the qualification of the interviewers in the HFCS, as their professional demeanor and expertise notably contribute to the quality of data. The chapter also covers details on the contact strategy and incentives for households to participate in the survey. Moreover, it summarizes the information material and documents that were made available to the households in the HFCS sample.

After the interviewers had collected the raw data, all information was reviewed already during the field phase; if necessary, further data were collected, or the data were Edited. This process is described in detail in chapter 4 on *Consistency Checks and Editing*. Each intervention in the raw data is thus presented in a transparent way, and additionally closed with flags in the data set documenting any alterations to the observations.

Chapter 5 on the *Multiple Imputations* applied in the HFCS deals with item nonresponse. For cases in which interviewees did not want or were unable to answer one or several questions, multiple imputations using statistical methods were applied 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 this method of dealing with missing observations. Like in the editing process, we use flags to indicate imputations. There-

fore, users are free to apply imputations or deal with item nonresponse in a different way.

A complex survey sample design was developed to ensure a sufficiently representative sample of households in Austria that fits the purpose of the Eurosystem and the OeNB. Chapter 6 on *Sampling* provides a detailed description of this survey sample design.

On this basis, the final household weights were calculated in several steps. Chapter 7 outlines the procedure of the *Construction of Survey Weights*. The sample 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 on the distribution of certain household characteristics).

Another step is required to obtain the correct variance estimation, which is presented in chapter 8 on the *Construction of Replicate Weights for Variance Estimation*.

The *User Guide* in chapter 9 provides basic advice on the correct use of HFCS data in Stata[®]. All other essential documentation and background material used in the HFCS in Austria is published in the *Online Appendix* at www.hfcs.at. This website will also provide information about the publication of HFCS data of all participating countries by the ECB in spring 2013 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. Before conducting the Household Finance and Consumption Survey (HFCS), it was possible to harmonize ex ante the information to be collected, and for the most part the methods used, despite cross-country differences in the technical conditions for its implementation, for example with respect to sample design (chapter 6) and multiple imputations (chapter 5).

The definitions of the variables collected in the HFCS were based on those used in three internationally renowned surveys: the Federal Reserve's *Survey of Consumer Finances* (SCF) (see Bricker et al., 2012), the Banco de España's *Survey of Household Finances* (EFF) (see Bover, 2011), and the Banca d'Italia's *Survey on Household Income and Wealth* (see Banca d'Italia, 2012). This chapter presents the Austrian questionnaire, which was elaborated on the basis of an English-language template.¹ The Austrian survey covers both the internationally agreed variables and country-specific characteristics (e.g. foreign currency loans or the national variant of housing association apartments).

This chapter is structured as follows: Section 2.2 outlines the objectives of the HFCS survey in Austria. Sections 2.3 and 2.4 focus on the definitions of the unit of collection and the reference period, respectively. Section 2.5 describes the sequence of the questionnaire and highlights some core questions and variables. Section 2.6 focuses on special features of the questionnaire, section 2.7 gives an overview of the documentation provided to interviewers, and section 2.8 lists the other participant countries. Finally, section 2.9 contains information about an online appendix.

2.2 Objectives of the Survey

The main goal of the HFCS was to collect microdata on the structure of euro area households' assets and liabilities. These survey data, which allow us to analyze households' investment and consumption decisions, 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 is the most comprehensive household-level survey conducted on this subject in the euro area. While a number of surveys had previously been carried out in individual European countries,² they were not harmonized, so that comparability of the results was limited. What is more, most of them did not cover all data required to sufficiently assess the households' overall wealth situation.

Especially for a central bank, household-level finance and expenditure data are indispensable, as they contribute significantly to improving the analysis of

¹ For further details about the network and the survey (including the English template questionnaire), see www.ecb.int/home/html/researcher_hfcn.en.html (accessed on January 22, 2013).

² See section 2.8 for a list of the participant countries and the respective web addresses.

monetary policy and financial stability. The most recent financial crisis has shown that the aggregate level of household debt – which can be calculated using macrodata – is not the key parameter for assessing stability risks, but that an analysis of the specific burden on different income, occupational and age groups is essential instead. In Austria, the HFCS was conducted with a view to closing this empirical gap and creating a better basis for monetary policy decisions and decisions concerning financial stability. The HFCS will also provide comparable structural data at the level of households in the euro area, which will also help improve forecast models and microsimulations.

2.3 Unit of Collection

For every survey, the target population (see also chapter 6) and the unit of collection have to be defined. While the main unit of collection in the HFCS is the household, some information is also collected at the personal level.

2.3.1 Definition of Household

For the purpose of the HFCS, a household is defined as one person living alone or a group of people living together in the same private dwelling who share household expenses and jointly take expenditure decisions. The following persons are to be regarded as household members:

- persons who live in the same household and are related to each other,
- persons who share household expenses and live in the same household but are not related,
- persons who usually live in the same household (observation period: past six months) 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 have their own household and are still financially dependent on the household.

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

Employees of residents, like au-pairs or nursing staff, short-time visitors or tenants are not considered household members. In shared apartments, all residents are treated as separate households unless they also share household expenses. This means that more than one household (e.g. people sharing a residence) can live at one address in line with the HFCS definition. 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 and households actually living in a residence that is registered only as a second home (or not at all) in the central population register.³

2.3.2 Financially Knowledgeable Person

The person who knows the most about household finances, i.e. liabilities, assets, income and expenditure, was selected by the household members to answer the

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

household-level questions. This financially knowledgeable person (FKP) answered all questions relating to the household as a whole (green sections in chart 1). The objective was to have all persons aged 16 and over to answer questions at the personal level personally. In case one or more household members were absent, the FKP could also provide personal information on behalf of these individuals.

While the FKP was in most cases a member of the household, this was not a prerequisite in 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.

2.4 Data Collection Period and Reference Period

In general, the questions asked in the HFCS referred to the status quo at the time the interview was conducted during the field phase from September 2010 until May 2011. In particular, this goes for all information obtained on stock data. In contrast, the questions about income (except the one about average monthly net household income)⁴ referred to the calendar year 2009 (reference period), as this was the last full calendar year before the survey was conducted.

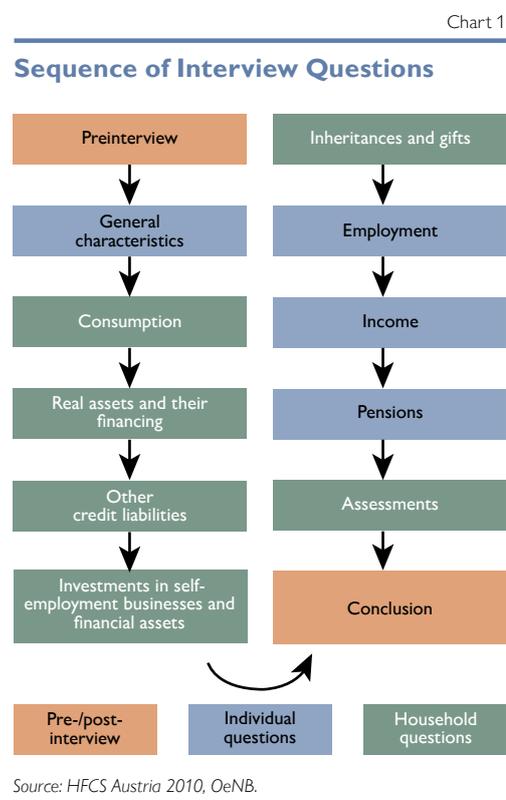
2.5 Interview Structure and Content

2.5.1 Questionnaire Structure

The questionnaire has three main parts: the preinterview, the main interview (divided into household and individual questions) and the postinterview. 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 an overview of the sequence of the sections of interview questions in the HFCS.

Preinterview

Before the interview, the household was informed about the content and structure of the survey.⁵ If the household was willing to participate, the interviewer recorded the household matrix and identified the FKP. Recording the household matrix data involved determining the size of the household and the total number of members in the household in line with the HFCS defini-



⁴ This is a non-core 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.

tion as well as identifying the FKP within or outside of the household. If the FKP was present, the interview was started immediately afterward – 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, information about the sociodemographic characteristics of all household members was collected: gender, age, migration status (if applicable). For household members aged 16 and over, information about the level of education (including that of their parents) and marital status was also recorded.

The subsequent sections (on consumption, real assets and their financing, other credit liabilities, 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 about some variables of total household income. The questions on consumption aim at collecting information about households' average monthly expenditure for food, total consumption and private transfers to persons outside the household. The statements on household income provided information used to assess whether this income was high enough to finance the household's expenditure, and if not, the interviewer asked how the household paid for any expenses in excess of income.

Real Assets and Their Financing

This section contained questions about the household's primary residence and most of the other real assets (excluding investments in self-employment businesses, which are addressed below) as well as their financing. For the sequence of the first half of the section, the ownership structure of the primary residence (variable (A)HB0300) played a central role. After the location and size of the primary residence had been determined, the households were asked whether they were (partial) owners, tenants or free users of their homes.

Homeowners were asked about the type of ownership transfer, the property value at the time of the interview and at the time of the ownership transfer as well as the rental value, about ownership of any real estate acquired before the main residence and about their attitude to precautionary saving. In addition, the interviewer asked whether the primary residence was financed through mortgages that use the residence as collateral. If so, the following data were collected separately for a maximum of three mortgage contracts: initial amount, duration of the loan, outstanding principal, repayment rates, interest rate and type of interest, borrower and other characteristics. If the household had more than three loans, the FKP was asked to provide the main information (outstanding principal and repayment rates) about all other loans taken together.

Households that lived in a rented home were asked about the rental cost (including and excluding operating expenses), the value of the property, about any property they may have acquired before their current residence, and about their

attitude to precautionary saving. In addition, tenants of housing association apartments were asked whether they had taken out a loan to finance the deposit made to the housing association.

Households that enjoyed free use of the primary residence were only asked about their attitude to precautionary saving, as none of the other questions in this section were relevant.

The next question for all three groups addressed additional real estate holdings. The interviewer also explicitly asked about holdings abroad. This section of the questionnaire sought basic information about the size, value and use of a maximum of three additional real estate properties (loop question⁶). For any properties in addition to these three, the interviewer collected only information about the total value. Like for the primary residence, the interviewer then asked separately about a maximum of three mortgage contracts that used these properties as collateral (loop question). If the household had more than three loans that used the additional properties as collateral, the interviewer again only collected summary information about all these loans.

Finally, the last questions in this section focused on the value of passenger cars and other vehicles as well as other real assets held by the household.

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, the interviewer asked whether the household held such liabilities and how much was outstanding (in the case of leasing contracts, the interviewer asked about 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 monthly payments, not about the monthly bill that is repaid at the end of a particular month.

The questions about uncollateralized loans were asked in a loop with up to three iterations. If a household held more than three loans of this type, the interviewer only recorded summary information about these additional loans (total outstanding principal and sum of repayment rates). In the loop, the respondent was asked about the same indicators as for the mortgage(s) backed by real property.

In addition, this section of the questionnaire addressed the household's risk attitude and the result of any application it might have made for a loan.

Investments in Self-Employment Businesses and Financial Investment

On the asset side of the households' balance sheet, all assets that were not already covered in the previous section on real assets were recorded. First, the interviewer asked whether the household owned part or all of any self-employment business that was not publicly traded, and, if so, recorded separate information on the industry, legal form, number of employees and value of up to three such businesses. Information about any additional investments in self-employment businesses was summarized to provide a total value.

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

The next questions focused on assets held in current accounts, savings accounts, saving and loan contracts, life insurance funds, mutual funds, bonds, listed stocks, silent partnerships, as well as assets in private foundations and in managed accounts. The FKP was asked for each of these components whether the household held such an investment (yes-no question) and, if the answer was yes, what the total value of these assets was. For all life insurance contracts, the interviewer asked about the date of conclusion and the total duration of the contract as well as the frequency and amount of payments into the life insurance plan, which allows the amount held in these funds to be extrapolated. In addition, the interviewer asked about financial debt owed to the household and other financial assets as well as the household's estimated total assets. 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, i.e. inheritances and gifts. In a loop, the interviewer asked about up to three inheritances or gifts, and recorded information about value at the date of transfer, the type of inheritance or gift, the source of the inheritance or gift, the year in which the transfer was made, and the recipient. The inheritances or gifts were listed in descending order by the value of the inheritance and/or gift for the current asset situation. For all inheritances or gifts in excess of three, only the value at the time of the transfer was recorded. The dataset also includes information about expected future inheritances or gifts.

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

Employment

The questionnaire collected information on the employment status of household members aged 16 and over. This section focused on the data of people active in the labor market, while pensioners, homemakers, students and unemployed people were only asked about their expected retirement age and about the number of years spent in employment so far. Then they continued with the next section of the questionnaire. Employed persons answered questions about their occupation (ISCO code⁸), the number of working hours (and seasonal fluctuations), the company's main economic activity (ÖNACE code⁹), the amount of time they had worked both for that particular employer and overall in their working life, and

⁷ For details, see chapter 4.

⁸ ISCO: International Standard Classification of Occupations, see www.ilo.org/public/english/bureau/stat/isco/index.htm (accessed on February 14, 2013).

⁹ 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). www.statistik.at/web_en/classifications/implementation_of_the_onace2008/index.html (accessed on February 14, 2013).

about their expected retirement age. The information provided in this section of the questionnaire is especially relevant in combination with that of the next section, which treats 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. All types of income were recorded for the calendar year 2009, as this was the last full calendar year before the start of the field phase in 2010, and information on annual income is usually readily available, e.g. on the income tax declaration.

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

Besides asking about income at the personal level, the interviewer also asked in the interview with the FKP at the household level whether the household received income from social transfers on a regular basis (both publicly provided social benefits and private transfers), rental or leasing income from real estate property, income from financial investments and income from a private business or partnership. Finally, the interviewer asked about other sources of income and about the expected income growth.

Pensions

The main pension variables collected in the HFCS questionnaire included the number of contribution years and the level of expected public pension benefits (as a percentage of income), as well as the number of contribution years, the payment rates and the balance in occupational or private pension accounts.¹⁰

Assessments¹¹

In this section, the FKP was invited to provide his or her assessment of social issues in Austria and of the household's position in the national income and wealth distribution.

Postinterview

After the interview, respondents were given the opportunity to record as a verbatim text any difficulties with specific questions, missing information, and other comments. In addition, this section includes information from the interviewer about the interview, the so-called paradata (see section 2.6.3).

¹⁰ As the respondents often found the questions on pension benefits difficult to answer, special caution must be exercised when evaluating the information in this section.

¹¹ This is a non-core variable specific to Austria that is not included in the international HFCS dataset.

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 PDF template, which is available in the online appendix. In addition to filtering, the questionnaire also contains internal consistency checks (chapter 4), which allow possible errors in data entry to be corrected already during the interview. The use of the CAPI (computer-assisted personal interview) 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 so that a specialized questionnaire could be provided for each household.

2.5.2.2 CAPI Testing

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

2.5.2.3 CAPI Problems

Despite comprehensive testing of the questionnaire, only a few difficulties persisted at the beginning and during the actual field phase:

- Category "Other" for country of birth (RA0400): While a verbatim option had been planned originally, it was not implemented. As a result, all respondents who chose the category "other" as their birth country – because their birth country was not in the list provided (see online appendix for the birth country list), e.g. because these countries had ceased to exist as such – were not routed to a verbatim option for the birth country but to the next question.
- Education of parents¹² (APA02\$0): This information should not have been collected for the FKP's children, as it had already been entered into the dataset with the FKP's response. This filter did not work fully correctly in all cases.
- Uncapped range for the value of investments in self-employment businesses (HD0801): At the beginning of the field phase, a programming error occurred in the euro loop for the question on a household's first investment in self-employment business. After indicating an open-ended range, respondents should have been routed to the currency information and the value confirmation step. Instead, they were asked to choose from a fixed range with an upper and a lower bound (one household was affected, see chapter 4).
- Variables of the section on employment (PE0900, PE1000, PE1100): The filter applied not only to the main employment status (as required by the ECB), but also to every other employment status. As a result of the error in the filter, the question was not included in the questionnaire for a few people who should have been routed to this question. The missing information was imputed for these respondents (chapter 5).

2.6 Special Features

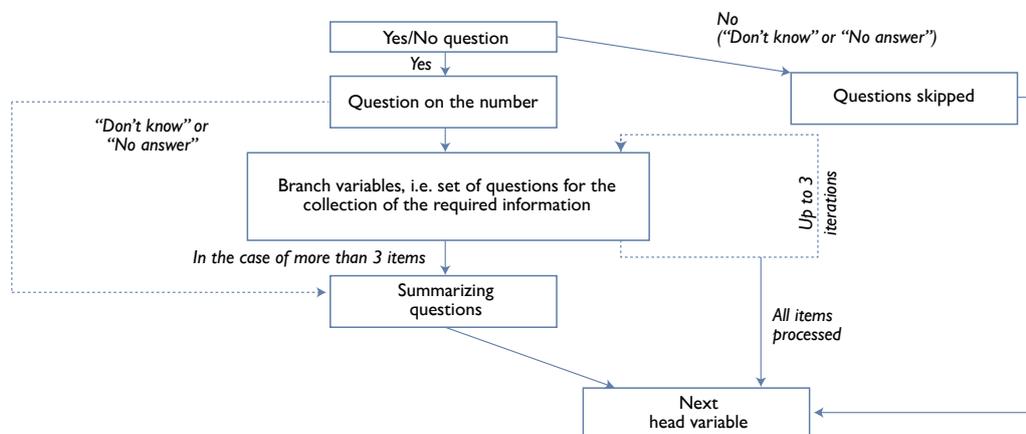
2.6.1 Loops

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 possible

¹² This is a non-core variable specific to Austria that is not included in the international HFCS dataset.

Chart 2

Structure of a Loop



Source: HFCS Austria 2010, OeNB.

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

Loops were used for the following items:

- mortgage loans that used 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,
- uncollateralized loans,
- holdings of investments in self-employment businesses, and
- ownership transfers through inheritances or gifts.

First, respondents indicated whether a specific item applied to their household, and if so, how many of such items the household holds. Then the details were recorded in a loop for every one of these items. For instance, if a household held two uncollateralized loans, the respondent was asked to first provide details on the loan with the higher outstanding balance and then on the second loan. This procedure was repeated up to three times. If a household held more than three uncollateralized loans, the interviewer recorded summary information about the ones in excess of three.¹³ The sequence of items was defined in the questionnaire.

To make the survey as user-friendly as possible, the interviewer could exit loops at any time. In the case of an exit, summary information was collected on the missing items. Interviewers were instructed not to overuse this feature; instead, it was only a means to prevent respondents from breaking off the interview early, if they were unwilling to answer the same 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 did not know, or want to state, the exact amount. This subsection describes the

¹³ Only very few households indicated that they held more than three of a particular item.

structure of such a loop. A schematic overview can be found in the document on euro loops in the online appendix.

In a first step, respondents were asked to provide exact amounts (“How much...?”). Respondents could state known amounts in any given currency. After the currency 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 exact amount was provided, respondents were asked to indicate a range with an upper and a lower bound. (“Could you provide a range with an upper and a lower bound for the amount?”) The range could 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 with an upper and/or lower bound, the interviewer continued as before, asking respondents to first indicate the currency and then confirm the range with upper and lower bounds 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 (see online appendix, euro loops document and catalogue of show cards). In this case, the only currency available was euro, and the interviewer moved on to the confirmation question as soon as the respondent had chosen a range. The same list of predefined ranges was used for all questions involving monetary amounts, so that the interviewer could use one and the same show card for all these questions.

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 recordings in ranges were especially important for the multiple imputations (chapter 5).

2.6.3 Additional Information

2.6.3.1 Contact Attempts

Every household in the sample population had to be unsuccessfully contacted on at least five separate occasions before it could be classified as a unit nonresponse (see also the contact rules in chapter 3).¹⁴

In the dataset additional to the collected information, these contact attempts were recorded in variables¹⁵ SB010\$ to SB0900. 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 to record in the electronic questionnaire once a household interview had been concluded.

2.6.3.2 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 category covered all information that could be obtained without actually entering a household’s residence or completing an interview: the inter-

¹⁴ Interviewers had to make at least two contact attempts in person over a period of more than two weeks.

¹⁵ Because anonymization rules had to be observed, these variables were not included in the user database.

viewer's assessment of the building and construction type, the geographical location (urban or rural area), the state of the building and special security measures.

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

The first type of paradata was recorded in the sample register file, which is not part of the user database because anonymization rules had to be observed. It was used mainly to calculate nonresponse weights.¹⁶ The second type was recorded in variables HR0100 to HR1600 in the household data file, which is part of the user database.

2.7 Documents Provided for the Interviewer

Among other things, interviewers had access to the following documentation to help them prepare for an interview and as a source of reference during the interview (also available in the online appendix):¹⁷

- a catalogue of show cards, which was used during interviews to provide respondents with a list of possible response options for several questions in the questionnaire,
- a glossary in alphabetical order and by subject, which contained simple definitions of the terminology used in the questionnaire, and
- a copy of the study entitled "Housing Wealth of Austrian Households" (Fessler et al., 2009) to illustrate how the data obtained might be used for analysis.

2.8 Participant Countries

The Eurosystem HFCS was conducted in the following euro area countries: Belgium, Germany¹⁸, Estonia (not included in the first wave), Ireland (not included in the first wave), Greece, Spain¹⁹, France, Italy²⁰, Cyprus²¹, Luxembourg, Malta²², Austria, the Netherlands, Portugal, Slovenia, Slovakia and Finland.

¹⁶ These weights are used for the correction of nonrandom participation of households in a survey and are needed to construct the final household weights (chapter 7).

¹⁷ For a detailed description of the documentation, see chapter 3.

¹⁸ Information on the survey in Germany (in German) can be found at www.bundesbank.de/Redaktion/DE/Standardartikel/Bundesbank/Forschungszentrum/studie.html (accessed on January 22, 2013).

¹⁹ Information about the survey in Spain (in Spanish) is available at www.bde.es/bde/es/areas/estadis/Otras_estadistic/Encuesta_Financi/ (accessed on January 22, 2013).

²⁰ Information on the survey in Italy (in English) is available from www.bancaditalia.it/statistiche/indcamp/bilfait;internal&action=_setlanguage.action?LANGUAGE=en (accessed on January 22, 2013).

²¹ Information on the survey in Cyprus (in English) is available from www.centralbank.gov.cy/nqcontent.cfm?a_id=11840&lang=en (accessed on January 22, 2013).

²² Information on the survey in Malta (in English) can be found at www.centralbankmalta.org/site/hfc_survey.html (accessed on January 22, 2013).

The survey was prepared by the Household Finance and Consumption Network (HFCN), which is organized by the ECB. The aim was to achieve ex ante harmonization at as many levels of the survey as possible. In doing so, it was necessary to take features specific to individual countries into account, leading to discrepancies and to additional questions in some cases. In addition to the core output variables, the Austrian survey also collected data that are specific to Austrian households (e.g. information on foreign currency loans). Moreover, the answer options for some questions were categorized in greater detail in the national datasets. A case in point is the question about the respondents' marital status, which had six answer options in the national dataset but only five answer options in the international dataset. The OeNB is planning to provide country-specific information in addition to the data sets that the ECB will publish in spring 2013.

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 document,
- the paradata questions,
- the country of birth list,
- the variables list,
- the catalogue of show cards, and
- the glossary.

3 Interviewers

This chapter provides an overview of the HFCS interviewers' role and tasks. It describes how interviewers were trained, supported and monitored 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 commonly considered sensitive. Therefore, the personal interviews conducted by trained interviewers played a major role in the HFCS process. Interviewers' professionalism, profound knowledge on 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 a comprehensive training on HFCS content and structure.

In the field phase and during the personal interviews interviewers had the option of consulting written reference material and, if necessary, of receiving support from the OeNB.

3.2 General Information

A total of 85 interviewers were deployed during the HFCS. While the survey company decided which interviewers to deploy for this complex and sensitive survey, the OeNB reserved the right to withdraw individual interviewers if they did not meet the quality criteria.

In general, interviewers had to have experience with surveys of this magnitude (for instance as a result of taking part in surveys such as the OeNB Household Survey on Housing Wealth 2008, EU-SILC or SHARE) to be deployed. More than 85% of the HFCS interviewers met this criterion. Payments for successfully completed interviews were calculated on the basis of the surveyed household's size; a considerably lower remuneration for the collection of paradata was paid in case interviews were not conducted successfully. Travel expenses were also refunded. To be entitled to a refund of travel expenses for unsuccessful interviews interviewers were required to have made at least two personal contact attempts and five contacts attempts altogether.

3.3 Interviewer Training

All interviewers conducting interviews as part of the HFCS were specially trained. Training was organized by the survey company and took place in cooperation with an OeNB expert from the HFCS team prior to the field phase at a total of five one-day training events in Vienna (two), Graz, Ansfelden and Landeck. Each training day comprised four units with the following content:

3.3.1 Training Unit 1

Introduction

First, the interviewers were introduced to the topic and the aims of the HFCS in Austria. The analytical approaches were explained on the basis of the results of a pilot study (OeNB Household Survey on Housing Wealth 2008). It was emphasized that it was particularly important that the interviews were conducted conscientiously and that each household in the sample participated in the survey.

Administrative Provisions

This session focused on the provisions for establishing contact, on interviewers' tasks, the reasons for households to take part in the survey as well as on issues of data quality and security. The interviewers were provided with contact details of contact persons at the survey company and at the OeNB in case they had any further questions.

Interviewer Behavior

In this session, issues concerning the behavior of the interviewers, particularly regarding further questions and inquiries of the respondents during the interview, were discussed and the interviewers were trained to refrain from expressing their own opinions during the interview. They were also instructed to repeat and explain questions as neutrally as possible (by using the glossary, if necessary).

3.3.2 Training Unit 2

Structure and Content of the Questionnaire

Based on the content covered in training unit 1, the structure and content of the questionnaire were outlined in detail (see chapter 2) as was the exact definition of the term "household" within the meaning of the HFCS.

Special Features of the Questionnaire

In this session, the interviewers were familiarized with the special features of the questionnaire such as the recording of liabilities in loops, loops for euro-related questions and paradata (see also chapter 2).

3.3.3 Training Unit 3

Test Interview

An integral part of the training consisted in conducting notebook-based test interviews. In accordance with OeNB provisions, the test interviews were conducted with a hypothetical household that comprised several members at an owner-occupied main residence which the owner had used as collateral to secure loans. Particularly difficult sections of the questionnaire, e.g. the recording of the value of business equity, were simulated. Each interviewer conducted an entire interview during this training unit.

3.3.4 Training Unit 4

Clarification of Technical Terms

In the last training unit, technical terms used in the HFCS were explained on the basis of the glossary (and with reference to the questions that had arisen during the training).

Interview Documents

All supporting documents available to the interviewers were discussed once again. Throughout the entire training, the interviewers were explicitly encouraged to ask questions and to raise issues that were unclear for every single aspect of the HFCS. This helped identify and avoid many potential difficulties prior to the start of the field phase.

3.3.5 Additional Training Sessions

Only trained interviewers were deployed in the HFCS. However, it was not possible to instruct all interviewers in the main training sessions as a result of terminations and new recruitments at the survey company. This meant that the survey company trained additional interviewers during the field phase. The scope of these additional training sessions was considerably smaller than that of the main training sessions, but so was the size of the groups with a maximum of three persons.¹

3.4 Contact Strategies and Contact Provisions

The process of establishing contact with the households in the HFCS sample took place according to detailed provisions by the OeNB. Between one and two weeks prior to the first personal contact attempt by the interviewer, the survey company sent the households selected in the sample a personalized advance letter signed by the OeNB Governor as well as an information folder. 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), 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 account statements, etc.) prior to the interviewer's first personal contact attempt.

Once the advance letters had been sent to the households, interviewers had to make up to five contact attempts with each of them. At least two of these contact attempts were to be made personally (by passing at the household's address in person and trying to establish contact); at least one attempt was to be made at a weekend and another outside normal working hours (9:00 a.m. to 5:00 p.m.). All contact attempts had to be spread out over a period no shorter than three weeks. This approach was necessary to rule out distortions caused by 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 interviewer was required to document every contact attempt. During at least one of the personal contact attempts, information on the exterior and the location of the respective property (see also section 2.6.3.2 on paradata) had to be 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, 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 it requested time to consider or wanted to make an appointment, or if it 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 the respondents outside their main residence (e.g. at the respondent's office).

¹ For the next wave of the HFCS in Austria, which is scheduled for 2014, these additional training sessions are to be replaced with additional training events that match the scope of the main training sessions and take place during the field phase in cooperation with an OeNB expert.

3.5 Documents and Other Supporting Material

In addition to being prepared in the training sessions for the interviews they were to conduct with the households, interviewers were also 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 a personalized letter and an information folder (see online appendix) explaining what the survey was about, what objective it served and how the collected data would be used. Households were assured that the collected data would be processed confidentially, and contact information was provided for further questions.

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 (corresponding to an actual value of around EUR 15 at the time of the field phase). The interviewers handed out 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 carry a copy of the study *Housing Wealth of Austrian Households* by Fessler et al. (2009) during each contact attempt. On the basis of data taken from the OeNB Household Survey on Housing Wealth 2008, this analysis gives an example of how survey data are used in a statistical context. This helped the respondents understand in advance how the information they provided was going to be used and thereby build up confidence in the survey. This study is also part of the documentation in the online appendix.

3.5.4 Catalogue of Show Cards

To answer certain questions of the survey, the respondents had to choose from a list of answers presented by the interviewer on show cards. These cards were part of an entire catalogue of show cards (see online appendix) which had been given to the interviewers beforehand and which covered the following topics:

- Intervals for euro amounts
- Methods to cover expenditure
- Types of income
- Rent including operating expenses
- Loan repayment
- Lending institution
- Main reason for renting
- Valuables
- Economic sectors
- Types of mutual funds
- Banks

- Investment behavior
- Type of inheritance/gift
- Employment status I and II
- Assessments

The questions that required the interviewer to provide a show card were specifically marked in the questionnaire (see online appendix). The digital version of the questionnaire also contained references in the relevant places to the appropriate card that was to be shown when asking the respective question.

3.5.5 Contact Form

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

Besides the household's identification number, documentation comprised the date, time, type (e.g. personal or by telephone) and outcome (e.g. complete interview or invalid 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 Handbook

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

3.5.7 Glossary

The content of the HFCS required a basic understanding of the different financial instruments, investment opportunities and types of income as well as the acquisition of nonfinancial assets. With this in mind, interviewers were given a glossary (see also online appendix) that provided explanations of technical terms both in alphabetical order and in order of the questionnaire's structure. On around 20 pages, the glossary included explanations for all terms of key importance to the HFCS, such as *investment funds* or *household* (within the meaning of the HFCS definition).

Already at the training stage, the interviewers were encouraged to use this glossary to acquire relevant knowledge which they would be able to fall back on

during the interviews. As the glossary refers to the variables recorded in the survey and explains the technical terms contained in the questionnaire, it is also of importance when analyzing the collected data.

3.6 Monitoring

To uphold the high quality standards of the HFCS, both the survey company and the OeNB monitored the interviewers. 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 particularly monitored the correct conduct of the interviews by checking on around every one in six interviewed households via telephone from Vienna. During these calls, the contacted respondents were asked to provide data on the composition of the respective household, the conduct and duration of the interview and the topics covered.

During the field phase, the data on the completed household interviews were forwarded to the OeNB promptly and in seven batches. Based on expert analyses of the household sheets, which contained household-level information from the interviews and were evaluated individually and grouped by interviewer, it was possible to monitor interviewers' performance. In addition, interviewer performance indicators such as item nonresponse (both broken down by nonfinancial and financial assets and in aggregate form for the entire interview), the relative duration² 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 were examined. The specific comments by the interviewers, which were required upon completion of each household interview, were examined as well.

At this stage, the OeNB aimed to quickly identify and resolve difficulties through timely analyses and, if necessary, to immediately withdraw those interviewers from the survey that did not meet the requirements.

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 the subsequent interviews. The next batch of data was then examined with particular attention to whether interviewers continued to make the same types of errors. In a few cases, misunderstandings were found in relation to the reference period for income (for most of the surveyed income variables, the gross income recorded referred to the 2009 calendar year) 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, the monitoring also helped reduce the item nonresponse rate of the households they interviewed.

Three interviewers had to be withdrawn from the survey rather soon after the field phase started due to flaws in conducting the interviews. In these cases, the

² During each interview, time logs were recorded at different points in the questionnaire.

household interviews they had conducted and which had yielded data fraught with considerable uncertainties were conducted once again by other interviewers or classified as unsuccessfully completed interviews after the field phase (17 households on the whole).

3.8 Online Appendix

The letter by the OeNB's Governor to the households, the information folder, the catalogue of show cards, the interviewer handbook, the alphabetical glossary as well as the study by Fessler et al. (2009) are available in German in the online appendix at www.hfcs.at.

4 Consistency Checks and Editing

4.1 Introduction

Data editing is understood to mean the ex post amendment of electronic records of 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 proxies, 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 of 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, interviewers may have recorded information incorrectly (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 HFCS 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 collected observations (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 22,000 of the close to 841,000 observations collected through the HFCS were edited, i.e. a mere 2.6% of all data points are amended (see the “All edits” column in table 1).

Table 1

Number and Type of Edits

	All ¹	Type of edits		
		based on expert analysis ²	based on proxies ³	deleted values ⁴
Total observations ⁵	840,714	840,714	840,714	840,714
Number of edits	21,837	6,867	13,767	1,203
Share of edited observations in total observations (%)	2.6	0.8	1.6	0.1

Source: HFCS Austria 2010, OeNB.

¹ All edits.

² Number or percentage share of edits made on the basis of expert assessments.

³ Number or percentage share of edits made on the basis of other survey information (e.g. verbatim records).

⁴ Number or percentage share of deleted observations.

⁵ Observations refer only to observed information. Observations that are filter missings are excluded.

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

The three columns next to “All edits” indicate the different types of edits. Editing resulted in changes to the collected values in the case of less than 6,900 observations. These changes involved primarily inconsistent values that were corrected through following investigations and/or other information or were deleted and imputed on the basis of the imputation model. Two-thirds of all of the amendments, i.e. about 13,800 observations, could be deduced from the verbatim records and by using the responses given to user-friendly questions on, say, respondents’ life assurance contracts or their total annual net income. All in all, some 1.6% of all observations were amended through edits of this kind. This low figure can be seen as an indication of the successful design of the questionnaire and the special training of the interviewers (chapters 2 and 3). In 1,203 cases, finally, i.e. for only 0.1% of all observations, the values collected were deleted and replaced with the filter missing (“.”) – among other things because a head variable was edited.

A case in point² would be the duplicate recording of income from pensions, first under “received employee income” and second under “received income from public pensions.” Here, the head variable “received employee income” (PG0100) was changed to “No” and the value (PG0110) entered under this income variable was deleted because the respective income figure had already been correctly recorded under the pension income variable (PG0300 and PG0310).

4.3 Checks for Consistency during the Interviews

The HFCS was based on Computer-Assisted Personal Interview (CAPI). The CAPI format has a number of advantages over the use of paper-based questionnaires or phone-based interviews. The interviewer used a laptop on which the survey software had been installed and was guided through the questionnaire on screen. The information collected was checked for appropriateness and consistency as it was being entered. Any questions of clarification that the respondents might have raised could be resolved immediately either by the interviewer or with the aid of the explanatory documentation at hand, so that errors could be avoided already at the time the data were recorded.

However, checks for consistency in the course of an interview are subject to limitations with regard to both substance and number. An excessive number of consistency checks during an interview would make it exceedingly long. The respondents are “worn down” and the standard of the data collected would decline. Interviews might even have to be broken off in individual cases.

Restrictions of substance arise from the fact that the details against which answers are meant to be checked for consistency need to have been collected in the first place. 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 right away. Yet the information necessary for more complex consistency checks often does not become available until the latter stages of the interview as it first needs to be generated through other sets of questions.

² Examples given in this chapter are indented for ease of reference.

The digitalized version of the questionnaire used for the HFCS provided for more than 150 consistency checks,³ typically in the form of “soft” checks: Whenever a test criterion was violated, a window popped up on the laptop screen to draw attention to the inconsistency of the answer given.

If a household with a disposable net monthly income of EUR 1,000 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 highlighted pop-up appeared in the questionnaire input mask:

“The sum of total consumption expenditure and regular remittances to nonhousehold members exceeds the household’s total net income. Are the figures correct? Please confirm that they are, or amend the figure(s) as necessary.”

As the figures collected may have referred to different moments in time, or the remittances might have been financed through sales of assets, or as the household’s income might have dropped as a result of one or more members losing their job, the correctness of the figures provided could not be ruled out. In such a case, the interviewer would ask the respondent to confirm or correct the household’s total income, its remittances and its consumption expenditure.

Other consistency checks programmed into the digitalized version of the HFCS questionnaire in Austria made it possible to proceed with the survey only if an answer that had been recognized as being incorrect or inconsistent was amended accordingly. However, these so-called “hard” checks were only used in cases where a particular answer could definitely be ruled out.

If a person stated that, say, he/she had lived in Austria for 40 years but gave his/her age as 30, the following highlighted 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 an alternation of both variables).

4.4 Postinterview Checks for Consistency

4.4.1 Analysis of the Data by Experts

During the fieldphase of the HFCS in Austria, the data on households deemed to be final by the survey company were forwarded to the OeNB in seven installments. Upon receipt, all data were analyzed by experts right away.⁴ 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.

³ A list of all the consistency checks that were programmed into the digitalized version of the questionnaire can be found in the online appendix.

⁴ The evaluation was carried out with the aid of external sources of data such as the OeNB’s 2008 Household Survey of Housing Wealth and the EU-SILC (conducted by Statistics Austria).

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 have created a bias towards certain households (e.g. housewife and pensioner bias) and distorted the data accordingly. The interviewers were aware of the fact that the list of addresses was limited to the 4,436 households of the gross sample. This made it possible to ensure that interviewers would not select the less difficult households and then move on to a new set of addresses. The incentive for interviewers to handle the strictly limited address material as efficiently as possible was supported through a merit pay 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 the latter had made an adequate effort to survey the households they had been assigned. The decision to exclude subsequent draws (substitute households) is among the key criteria for a successful survey, and moreover essential for ensuring the representativeness of the sample (Vehovar, 1999).

Initial analysis of the information on individual households during the field phase covered the information 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 interviewers, as well as the date, time and duration of the interviews. This set of information enabled an initial assessment of the quality of the interview. The microdata on every single household were checked for consistency of substance and reviewed by at least two analysts. 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 draw their attention to errors or incomprehensibilities. The shortcomings identified in this process were often of a minor nature, but three interviewers whose results were not up to the required standards (e.g. regarding nonresponse) were excluded.

4.4.2 Follow-Up Investigations

If the data analysis did not lead to the root of a data problem, households were contacted again by the survey company to clarify uncertainties and ensure that data were recorded correctly. A typical case of a data problem that was easy to spot and did not require follow-up investigations was rewriting a negative current account balance as a positive liability (account overdraft) while setting the current account balance to zero (see also section 4.6). This was simply a matter of adhering to the recording conventions for such liabilities. Decisions on follow-up investigations were always guided by the principle that any ex post editing of data and additional burdens on participating households should be kept to a minimum. All in all, a follow-up was necessary on specific details of some 400 households. Many of the unusual results (e.g. particularly high amounts of assets) were confirmed or else corrected in the course of the follow-up investigations.

4.4.3 Investigation of Outliers

Particular attention was paid within the scope of the individual analyses to the recognition and processing of outliers (exceptionally high or low values), which were recorded above all for financial variables, the size the household income or the size of the dwelling. Any outliers that were not removed from the dataset were actually not the result of interview errors but largely confirmed in the follow-up. It therefore seems appropriate in future studies based on HFCS data 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 Filters and Consistency

On top of the consistency checks programed into the digitalized version of the questionnaire and the analysis of the data by experts, the field phase also included detailed automated consistency checks of the data of all households.

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. Wherever the programming of individual hard checks was found to be faulty, the polling firm was informed so that the fault could be corrected.

The technical review also covered the questionnaire's complete set of filters to prevent programming errors from leading to extensive and costly follow-up interviews. Comprehensive tests of the questionnaire's programming prior to the start of the field phase as well as a pilot survey of 50 households led to the identification and correction of minor programming errors. For instance, whenever two members of a household refused to reveal their age, the second individual's answers to most of the personal questions were suppressed initially (see section 4.6.2.12).⁵ These filter checks also made it possible to ensure that the coding of variables was consistent throughout the questionnaire.⁶

4.5 Flags

All edits (and imputations – see chapter 5) can be identified with flag variables, which document in detail how the individual HFCS observations were established (see table 2 for a list of the flags used to classify the observations). To comply with international provisions, we aggregated a number of flags to ensure that the datasets are comparable at the international level (section 4.7).

Group I

The flags allocated to group I were used to indicate that the data points had been compiled. Specifically, all the values documented with the survey software during the interview were assigned flag 1, while all filter missing observations (".") were flagged with a 0. Information recorded in loops (section 4.6.2.4) was – where necessary – moved in the iteration of a loop and assigned flag 2. This means that flag 2 observations are included in the dataset exactly as they were recorded but carry a different iteration number.

⁵ This problem was discussed and resolved soon after the receipt of the first installment of the survey data, so that it remains limited to only a few households.

⁶ All HFCS variables were assigned value labels that explain the coding. The coding of the individual variables is also included in the questionnaire (available in the online appendix).

Table 2

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
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 previous question
	1053	Not imputed, originally collected from a range or from brackets
	1054	Not imputed, collected value deleted or value not collected due to a CAPI or interviewer error
	1055	Not imputed, edited to "missing" due to incorrect answer to a higher-order question
Group III	2050	Missing, set as "missing" for anonymization purposes
	2051	Missing, set as "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 as "missing" (":")
Group V	4050	Imputed, originally "Don't know"
	4051	Imputed, originally "No answer"
	4052	Imputed, originally not collected due to missing answer to a previous question
	4053	Imputed, originally collected from a range or from brackets
	4054	Imputed, collected value deleted or value not collected due to a CAPI or interviewer error
	4055	Imputed, originally not recorded due to incorrect answer to a higher-order question

Source: HFCS Austria 2010, 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 an exact 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 1055). Observations with group II flags were not imputed (chapter 5).

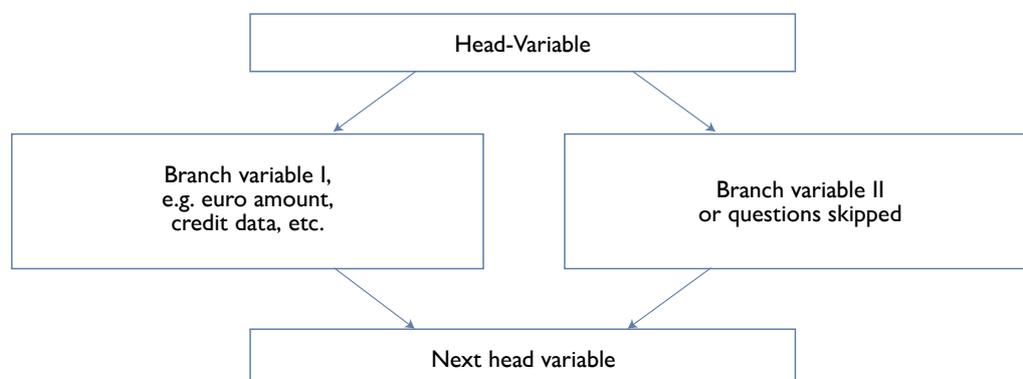
Group III

Observations and/or variables that were not recorded or that were recorded but subsequently deleted from the datasets on account of anonymization rules were assigned group III flags.

Group IV

Flags belonging to group IV indicate ex post edits of an observed value. The following types of ex post edits can be distinguished: edits required by logical inconsistencies were flagged with 3050; observations that were adjusted using other information obtained in the survey (e.g. that on life assurance contracts – see section 4.6.2.9 for details) were flagged with 3051; observations that were adjusted on the basis of verbatim records (see section 4.6.2.3) were encoded with 3052; and edits to delete a value and set the observation set to filter missing (as in the case of duplicate entries) were flagged with 3053. All observations that were corrected through follow-up investigations were assigned flag 3050.

Sequence of Questions



Source: HFCS Austria 2010, OeNB.

Group V

Flags in group V mirror those of 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 rather than an exact figure, which was subsequently imputed, this observations was flagged with 4053 after multiple imputations. This ensures that all information can be tracked even after imputations.

Chart 3 indicates how questions were typically structured in the HFCS questionnaire. Let us take employee income to give an example of the tree of questions⁷ 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 (to be answered with “Yes” or “No”). If the question was answered with “Yes,” the amount was recorded and the interview continued with the next head variable in the questionnaire – in this case, the question on self-employment income. If a household had no income of this kind, or if the respondent failed to provide the necessary information (answering the question with “Do not 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 with with 1 or 0. If it was found later (e.g. as a result of information on the employment situation in the household) that the question on employee income had been incorrectly answered with “No,” the initial response was corrected and flagged as “Edited, set to modified value as considered incorrect or unreliable” (3050) and the corresponding field for entering the amount was released for imputation. Following imputation, the value was then reflagged as “Imputed, not recorded due to incorrect answer of a higher-order question” (4055).

If, for instance, the question on the educational attainment of a member of the household (variable (A)PA0200) had been answered by selecting the category “Other qualifications” and if that answer was subsequently found

⁷ See chapter 2 for details on the structure of the questionnaire.

to match one of the predefined categories, the observation was flagged as “Edited, adjusted on the basis of the verbatim records” (3052) in the flag variables of the personal dataset.

This flag system allows the origins of every single observation in the HFCS to be tracked. No flags were used to encode the variables for identifying households and persons, nor were the country codes and the iteration number of the imputation flagged. The flags described here provide for a more detailed breakdown by category than those incorporated in the international HFCS dataset that can be obtained from the ECB. For reasons of international consistency, the flags were aggregated prior to submission to the ECB (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 investigations and ex post editing. Respondents’ answers were checked for plausibility against known benchmarks, as derived from descriptive statistics (e.g. on the average income) compiled on the basis of already completed HFCS interviews and external sources of data. Moreover, the review process heavily relied on proxies that recorded values in aggregated form and/or in a variety of other ways.

Both interviewers (chapter 3) that produced nonstandard results and follow-up investigations by the survey company were reviewed with particularly great care. Analyses by experts were generally used to resolve the following issues through ex post edits:

- Double entries: Cases where an inheritance, for instance, was recorded both under “Household main residence inherited” and in the “Inheritances and gifts” 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.
- Real estate ownership patterns in the context of the intra-household distribution of this asset item: A few cases where the household’s financially knowledgeable representative had taken the question on the ownership of the main residence to relate to himself/herself as an individual, rather than – as intended – to the household as such, had to be corrected.
- Implausible values: Values that remained implausible after follow-up investigations had to be changed into missing data entries that were subsequently imputed.
- Erroneous entries by interviewers: For example, the value 10 (for October 2011) that had been entered as the contact month for an interview conducted in 2011, the outcome of which had been submitted in February, had to be changed to read 1 (for January 2011) because the interview phase had ended in June.

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 investigations and/or the use of proxies (such as verbatim records) failed to provide further information – inconsistent observations were changed to missing and released for imputation. Inconsistent or implausible observations were processed with great care and only deleted if there was virtually no doubt about the inconsistency.

4.6.2 Structural Editing

4.6.2.1 Cleaning

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

4.6.2.2 Currency Conversion

Respondents could give amounts in different currencies (chapter 2). The edits set out below relate both to entries of exact amounts and to ranges indicated by the respondents (entries for predefined ranges had to be 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 13.7603.⁸ In addition, some amounts were given in Deutsche mark. In such cases, too, the irrevocably fixed conversion rate set by the ECB was used for conversion into euro, namely EUR 1 = DEM 1.95583.⁸

In the case of 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 of the exchange rates recorded in 2010, while the total 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 website⁹ being used as reference values.

4.6.2.3 Verbatim Records

For many questions, respondents were given the option of choosing the category “Other” and providing a verbatim response, essentially with a view to making the questionnaire as user-friendly as possible. Thus, recording of the verbatim description was allowed 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, as proved to be possible in most cases. Wherever this could not be done, the initial categorization of the variable as “Other” was retained. All observations subjected to ex post edits on the basis of verbatim records were flagged with flag 3052 (see section 4.5 for details on the flags used).

4.6.2.4 Navigation of Loops

As outlined in greater detail in section 2.6.1, some pieces of information were recorded in loops, which means that the interviewer ran through an identical set of questions for a sequence of items owned by the households. Information on the following items was collected in the form of loops:

⁸ For the irrevocably fixed euro conversion rates, see www.oenb.at/isaweb/report.do?lang=EN&report=2.12 (accessed on January 22, 2013).

⁹ For the effective exchange rate indices of the euro, see www.oenb.at/isaweb/report.do?lang=EN&report=2.16 (accessed on January 22, 2013).

- Mortgages on the main residence
- Real estate assets apart from the main residence
- Mortgages on such other property
- Unsecured loans
- Businesses owned by the household
- Inheritances and gifts

In the following 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 backed by the primary 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 had been asked to go through the loops of questions on inherited wealth in descending order of relevance for the household’s current wealth situation. At the same time, they were instructed to indicate only amounts as transferred rather than current amounts. Over the period between the time an item was inherited and the interview, certain inheritances can gain (or lose) more in value than others, or inherited residential property, for instance, might already have been passed on to children, so that it no longer has an impact on the household’s wealth situation.

Regarding flags, every variable within a loop that was replaced with observations recorded for the same variable in another loop was flagged with 2 (section 4.5). Wherever a filter missing for one of the variables within a loop was replaced with the filter missing for the same variable in another loop, flag 0 “Not applicable; skipped due to rerouting” was used.

Skipping of 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 with 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 (up to three) 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 of 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, inheritances and gifts, the relevant summary responses had to be edited accordingly. For ease of reference, edits are described in the following on the basis of the section of the questionnaire dealing with unsecured loans (section 2.5).

In the 21 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 indicated the outstanding amount (i) only in answer to the summary question, or (ii) both when running through

the first loop of questions and in answering the summary question, or (iii) neither during the first loop of questions nor in answer to the summary question. If the respondent had indicated the outstanding amount only in answer to the summary question – possibility (i) – this amount was entered as the answer to the appropriate question (in the first loop) and the entry under the summary question was edited to filter missing. If the respondent had indicated identical amounts – possibility (ii) – in answer to the loop question and under the summary question, the latter was edited to filter missing since it was a duplicate entry.¹⁰ Where no amount was given at all, neither within the loop nor in the summary – possibility (iii) – solely the summary question was edited to filter missing.

In cases where a household had taken out two unsecured loans, and had skipped questions within a loop,¹¹ the type of edit depended on whether the respondent had (i) specified the highest loan outstanding and indicated an aggregate amount in answer to the summary question; or (ii) indicated outstanding loans in answer to both loops of questions and the summary question; or (iii) specified amount solely in the answer to the summary question; or (iv) given no amounts at all, neither in the answers to the loop item questions nor in the answer to the summary question. In the case of possibility (i), the amount outstanding of 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 under the first loop, with that amount then being entered accordingly. This was, however, done only if the sum total of the two loans outstanding exceeded the amount outstanding of 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 loans outstanding, but rather the amount outstanding of the second such loan. In both instances, the summary question was subsequently edited to filter missing. In the case of possibility (ii), the amount given in answer to the summary question was edited to filter missing. If only the sum total of the two loans outstanding was given – possibility (iii) – it was used as the upper bound of both the first and the second such loan for the imputation model. If no amounts were given at all, neither under the loop questions for each of the two loans nor in answer to the summary question – possibility (iv) – the summary question was edited to filter missing.

The editing procedure followed in cases of 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. All edits were again properly flagged.

¹⁰ Where the amounts given were not identical, that given under the loop questions on the first loan was deemed to be more relevant than that given under the summary question. The reasoning behind this procedure is that the loop questions relating to the first loan contained a question asking expressly for the amount outstanding on an unsecured loan, so the amount given there is regarded as more trustworthy.

¹¹ This occurred in the case of only two households.

Summary Questions

Every loop of questions ended with summary questions (chart 2 in chapter 2). As a rule, the variables on summary questions contained in the dataset reflect only the residuals beyond any three items of a household. As indicated in chart 2, the summary questions were ultimately also put to all respondents who had refused to indicate the numbers of a given item in the household. In such cases of non-response, the information provided here was used for multiple imputations (chapter 5) and deleted from the dataset *ex post*.

4.6.2.5 Personal Variables at the Household Level

Various variables recorded information on the individuals belonging to the household, but they are stored in the household file. Examples of such information are details on the real estate ownership patterns within the household, on which member of the household took out which of the different loans or on that member of the household who works in the business owned by the household.

In order to be able to cover even unusually large households, variables were created for up to 18 individuals per household. However, the largest household interviewed successfully in Austria had only 9 members, so that all such variables in excess of that number were deleted from the dataset. Each such variable was reviewed individually, and those that had not been filled were deleted. If, say, individual No. 6 was the last member of the household in the list of borrowers to take out a loan, all variables for individual No. 8 and beyond were deleted. The variable for individual No. 7, while not containing a “true” entry in the sense that all households are encoded “No further individuals listed,” was retained in the dataset simply to reflect the latter.

4.6.2.6 Current Account Balances and Overdrafts

A number of households misreported a negative balance on their household current account as a current account balance (HD1110). There were also occasional duplicate entries, as well as misplaced entries, in this area that subsequently had to be edited.

4.6.2.7 Rent Variables

The HFCS questionnaire included questions on the amount of housing rent paid both excluding and including the running costs. In the case of some households, the rent excluding running costs was higher than, or equal to, that including such costs, which is simply not possible because housing cannot be “run” free of charge. Some of these households had entered only the running costs under the item “Rent including running costs.” In the course of editing, these were added to the amount entered under “Rent excluding running costs” to arrive at the “Rent including running costs.” In the case of other households, the “Rent including running costs” was edited to read missing and released for imputation, with the “Rent excluding running costs” serving as the lower bound to the “Rent including running costs.”

In addition, the item “Rent including running costs” was edited to become the upper bound for the variable “Rent excluding running costs” and used for imputations whenever the answer to the latter was not an amount (i.e. read “Don’t know” or “No answer”).

4.6.2.8 Agricultural Businesses

As defined in the HFCS, farmers are owners of an agricultural business. Some farmers, however, did not regard themselves as persons running a business. Such cases thus had to be reviewed separately and the observations had to be edited accordingly. The edits were based on the definition used to classify a household that owned an agricultural business. The classification was undertaken with the aid of the employment variables of all household members. In cases where at least one member of the household indicated that he/she worked (on a self-employed basis) as a farmer, the number of investments in self-employment businesses was increased by one only if no investment in a self-employment business had previously been recorded in the agricultural sector. If the investment in a self-employment business had already been recorded earlier, the entry was not edited. The NACE code for this business was set to that for agricultural businesses, and at least that member of the household who had stated that he/she worked as a farmer was deemed to be employed in the agricultural business. The legal form of the business was edited to read sole proprietorship. Both the household's ownership share and the value of the agricultural business were released for imputation.

In some cases, the value of the agricultural business was wrongly entered under value of the main residence. Account was taken of such information in imputations through delimitations and/or the use of a proxy variable, in which the value of the main residence was recorded together with the value of the investment in a self-employment business (section 5.3.5).

The category of agricultural businesses was subjected to special individual reviews by experts. Particularly complex cases were covered by follow-up investigations and logical amendments were made where necessary.

4.6.2.9 Life Assurance Policies

Information on assets held in the form of life assurance policies was recorded through questions ensuring that the answers were both as precise as possible and prone to only few errors. In particular, there was no direct question on the value of the assets in life assurance policies, but rather a series of questions on the start of payments, the frequency of payments (monthly or yearly) and the amount of the current payments for every single life assurance policy in the household.¹² The value of the assets held in the form of life assurance policies was calculated as the sum total of all payments. In cases where one or several details were not given, the remaining observations were used for the bounds of the value to be imputed.

4.6.2.10 Income Variables

The following categories of personal income were recorded separately for every member of the household who was 16 years old or older:

- Employee income (PG0110)
- Income from self-employment (PG0210)
- Income from public pensions (PG0310)
- Income from private and occupational pension plans (PG0410)
- Income from unemployment benefits (PG0510)

¹² Possible lump-sum payments at the start of a life assurance policy could be identified as such in the verbatim records.

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

- Income from public assistance or welfare payments (HG0110)
- Income from private transfers (HG0210)
- Rental income from real estate assets (HG0310)
- Income from financial investment (HG0410)
- Income from investments in self-employment businesses or partnerships (HG0510)
- Income from other sources (HG0610)

In the case of the first four personal income categories, respondents could indicate their net income if their gross annual income was not known (chapter 2).

Where only a net amount was entered for an income category, the gross income was calculated with the aid of the Austrian Federal Ministry of Finance's converter of gross to net income,¹³ with 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 children (if any), the Federal State and the employment status (white or blue-collar workers, pensioners)¹⁴ being used as a basis.

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 (to the extent that the legal requirements were fulfilled and the partner did not earn more than EUR 6,000 per annum).

Given the far greater scope that the self-employed have for tax deductions, a precise conversion was dispensed with in the case of income from self-employment. A precise conversion of the net into the gross amount is given only in the case of annual incomes of less than EUR 11,000, which are regarded as tax-free, so that the gross is equal to the net. For all other values (in the case of some 25 individuals), the net income was converted (on the basis of the employment status of white-collar workers), with EUR 10,000 subsequently being added to, or deducted from, the amount obtained in order to arrive at an interval for the imputation of the exact amount. This reflected the uncertainty that such a conversion entails, without losing the important information on the actual range within which the value is to be found.

If the net amount had moreover also been recorded only as a range, the upper and lower bounds were converted into gross values that were subsequently used in the imputations. All converted values were assigned flag 3051.

Income from private pension plans (PG0410) was not converted, but simply taken as gross amounts since the amounts involved were insignificant.¹⁵

Using flags as a basis, table 3 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 2,166 individuals. 1,263 of the respondents (58.3%)

¹³ See www.bmf.gv.at/service/anwend/steuerberech/bruttonetto/_start.htm (accessed on January 22, 2013, in German).

¹⁴ "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.

¹⁵ Seven individuals in all gave their income from private pension plans in net terms. For six of them, the annual income was below EUR 2,900, while one indicated a net income of EUR 18,500 per annum from private pension plans. Given their low value, all amounts were taken as gross amounts, so that there was no conversion.

Table 3

Number and Share of Edits of Gross Employee Income Based on Flags

	Number of persons	Share in %
Number of persons receiving employee income	2,166	100
Answer recorded, observation complete (flag 1)	1,263	58.3
Not imputed, originally "Don't know" (flag 1050)	49	2.3
Not imputed, originally "No answer" (flag 1051)	76	3.5
Not imputed, originally not collected due to missing answer to a previous question (flag 1052)	14	0.7
Not imputed, originally collected from a range or from brackets (flag 1053)	248	11.5
Not imputed, collected value deleted or value not collected due to a CAPI or interviewer error (flag 1054)	5	0.2
Not imputed, edited to "missing" due to incorrect answer to a higher-order question (flag 1055)	207	9.6
Edited, set to modified value as considered incorrect or unreliable (flag 3050)	28	1.3
Edited, adjusted on the basis of other information obtained in the (national) survey (flag 3051)	274	12.7
Edited, adjusted on the basis of the verbatim records (flag 3052)	2	0.1

Source: HFCS Austria 2010, OeNB.

expressed their annual income in gross terms, as had been requested. A further 49 respondents (2.3%) answered "Don't know" and 76 individuals (3.5%) opted for "No answer". Yet another 14 individuals answered the yes/no question on whether they had any employee income with either "Don't know" or "No answer," so that the flag (1052) precluded that they were asked about the amount. 248 respondents (around 11.5%) gave the amount of their income in the form of a range. The responses of 212 individuals (9.8%) were edited and flagged as missing and "to be imputed"; by far most of the respective edits (those of 207 individuals) were due to an incorrect head variable (flag 1055). 274 of the respondents (12.7%) were only able to provide their net income, which was then converted with the aid of the Federal Ministry of Finance's converter of gross to net income. The responses of the remaining 30 individuals (1.4%) were edited in value on the basis of analyzes by experts, with most edits due to follow-up investigations or to the aforementioned problem of additional or missing "zeros."

4.6.2.11 ISCO Classification

In accordance with the requirements of the ECB questionnaire, the main occupations of respondents were recorded (in variable PE0300) on the basis of the categories set out in the *International Standard Classification of Occupations* (ISCO88). Making individual members of each household classify their respective occupation themselves, however, would have been extremely difficult for respondents who do not have any advance knowledge of the ISCO categories, so that the observations received would have been inadequate. On those grounds, the questions on the job title and/or the tasks of relevance to the main occupation were asked orally in Austria, with the answers being recorded verbatim. That information was then transposed ex post into the corresponding ISCO code, as taken from Statistics Austria's list of Austrian ISCO designations ("Alphabetikum zur OEISCO"). As required by the ECB, classification was based on the ISCO aggregation of occupational groups (two-character codes). To this end, the verbatim record of the job title and related main tasks was supplemented with personal data of relevance for

the ISCO classification (in particular, the respondent's professional qualifications and the main activities of the company in which the respondent worked).

In order to ensure the international comparability of this data across all euro area countries, the ISCO codes used for classification were those of the old ISCO88, as required by the ECB. In the flags, the variable PE0300 to be submitted to the ECB was first flagged with 3051 – which means “Edited, adjusted on the basis of other information obtained in the (national) survey” – and aggregated in a next step (section 4.7).

4.6.2.12 Exclusion of Successful Interviews

For various reasons, the final data do not include those on the following three groups of, all in all, 28 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. boarding schools, homes for elderly people, prisons, etc.). A number of the households that had been interviewed successfully were excluded from the survey because the respondents were living a home for elderly people. The four households in this group were edited out and flagged as “Not belonging to the gross sample.” Two households that were not interviewed successfully were likewise eliminated from the gross sample and given the same status.
- Households with a high proportion of nonresponse items: This group, which comprised 22 households, had to be deleted from the dataset because the respondents refused to answer too many questions.
- CAPI errors: In the case of two households, a programming error in the digitalized questionnaire made it impossible to access the personal questions for all members of the household.

The observations for the households in the two latter groups were edited to “Nonparticipation on other grounds” and assigned a nonresponse weight of zero (section 7.2.3).

4.6.2.13 CAPI Errors Encountered with the Questionnaires

Category “Other” in Answer of the Question on the Country of Birth (RA0400)

A verbatim record was not possible for this category, which meant that “Other countries” could not be specified ex post and had to remain “Other” in the dataset. All in all, 11 individuals chose this category, and these answers were retained as such in the dataset.

Educational Attainment of Parents (APA02\$0¹⁶)

This question was not to be put to children of the household's financially knowledgeable person because the information on educational attainment had already been obtained from the latter. However, the filter did not work properly in all cases. Four individuals – parents of households' financially knowledgeable representatives – were not asked this question at all. The variables for these individuals was set to missing and released for imputation. In the case of a further 18 individuals who were all children of households' financially knowledgeable representatives, this question was asked by mistake, so that the data were deleted ex post.

¹⁶ This is a noncore variable specific to Austria that is not enclosed for in the international HFCS dataset.

Uncapped Range in the Case of Investments in Self-Employment Businesses (HD0801)

At the start of the field phase, the filter programming of the euro loop for the value of the first investment in a self-employment business was faulty. In cases where a range without an upper bound was entered, the filter did not direct to the recording variable the currency and the required confirmation, but rather to the variable for predefined ranges. This occurred only once, and the entry involved was obviously a euro amount and was also recorded as such.

Status of the Last Main Job and Duration of Employment (PE0900 and PE1000)

Faulty CAPI filter programming resulted in these variables being missing in the case of some individuals, while they were recorded unnecessarily in others. Accordingly, variable PE0900 entries for individuals who had been covered by mistake (i.e. for all individuals who were gainfully employed) were deleted in the course of the editing process in order to ensure conformity with international requirements. However, the information involved is easy to reinstate via the data recorded in answer to the next variable in the questionnaire (PE1000). In addition, the observations relating to 14 individuals who had mistakenly not been asked the question on the duration of their employment were released for imputation.

Expected Age of Retirement (PE1100)

The CAPI filter programming for the HFCS in Austria did not coincide with that of the ECB questionnaire in that the filter for the question on the age of retirement did not relate (as required by the ECB) to only the main job (PE0100a), but also to other jobs (HFCS conducted in Austria). For this reason, the question was not asked in the case of nine individuals. In order to ensure conformity with international requirements, personal observations under this variable were edited to missing and released for imputation.

Regular Payments of Occupational Pension Plan Benefits (PF0800)

The filter programming for the question on the payment of regular occupational pension plan benefits caused some individuals in Austria who had a claim to future payments of benefits under an occupational pension plan not be asked this question. The information on individuals who had mistakenly not been asked this question was imputed.

4.7 Formating and Editing after Multiple Imputations

Any information collected at a greater degree of granularity in Austria than in other countries 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 assigned categories under the International Standard Classification of Education (ISCED). As allocating domestic education paths to the respective ISCED is not straightforward, we recommend using the classification established for Austria.
- Labor/employment status: More detailed subcategories were aggregated.

- Main residence – ownership: More detailed subcategories were aggregated.
- Loan installments: 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 for Austria.
- 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 recorded in the category “Other.”
- Legal form of the business: More detailed subcategories were aggregated.
- Balances on savings plans with savings and loan banks and life assurance contracts: Data recorded on those two investment methods are aggregated into savings (HD1200 and HD1210).
- Investment behavior – willingness to take risks: The optional answer “No uniform allocation possible” was coded as “Don’t know.”
- Kind of assets received (survey questions on inheritances and gifts): The sequence based on asset values was dissolved.
- Provider of assets (survey questions on inheritances and gifts): More detailed subcategories were aggregated.
- Purpose of saving: The sequence by relevance was dissolved.
- Paradata: The variables HR1100 and HR1200 were recoded from a single-response to a multiple-choice answer and vice versa, while the nonrecorded variable HR14001 was edited to “missing.”
- Flags: The more detailed flags specific to Austria were aggregated to conform to international standards, i.e.
 - flag 1055 was recoded as 1054;
 - flags 3051, 3052 and 2 were recoded as 1;
 - flag 3053 was recoded as 0;
 - flag 4055 was recoded as 4054.

The additional data over and above those in 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 2013. The additional information includes additional variables, as well as a more detailed breakdown of certain variables. Datasets may be merged on the basis of both 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 most probably not been recorded correctly. In cases of ambiguity, the possibility of conducting ex post investigations on the phone was always considered first. Recourse to this option allowed many observations either to be corrected or to be confirmed as correct.

Knowledge of the steps undertaken to check the consistency of the data is essential both for any analysis of the data and for understanding how the observations have come about. 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 some other way.

The online appendix which supplements the information provided here on the edits and consistency checks applied in the HFCS in Austria contains a list of the consistency checks programmed into the digitalized version of the questionnaire.¹⁷

¹⁷ *All documents included in the online appendix are available at www.hfcs.at.*

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 is disregarded, this leads to biased estimates. For the HFCS data, we therefore used a method that addresses this problem, i.e. multiple imputation by 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 or, in other words, to maintain 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. In order to capture the uncertainty behind the missing values, we do not impute just one value per missing value, but several (in the case of the HFCS, five).

Other similar surveys – such as the *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 already imputed datasets. 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 opportunity 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 then describes the imputation procedure used, and in section 5.4 we explain the specification of the imputation model and how the imputations were executed. Section 5.5 finally presents some imputation results.

5.2 Item Nonresponse

Table 4 shows selected statistics on item nonresponse. On average, each household has 17.3 missing values, which means that item nonresponse was limited to a mere 2% of all the questions (variables) addressed to each household. However, this rate increases to 6.9% in the case of euro variables, which 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 implies that all households that have missing values in any of the variables of interest are deleted and that analyses are based on complete observations only. However, the loss of information resulting from this

¹ Another related problem that occurs in surveys is unit nonresponse, which means that no questions are answered at all because, for example, participation in the survey is declined by a household. We address this problem through the construction of HFCS nonresponse weights (chapter 7).

² For a comprehensive study, see Little and Rubin (2002).

Table 4

Item Nonresponse per Household (Unweighted)

	Mean	Median	Minimum	Maximum
Number of variables asked				
all variables	826.8	824.0	637	1,242
euro variables	52.1	53.0	17	98
Number of variables with missing values				
all variables	17.3	8.0	0	474
euro variables	3.6	2.0	0	54
Share of variables with missing values in %				
all variables	2.0	1.0	0	39.5
euro variables	6.9	4.2	0	78.8

Source: HFCS Austria 2010, 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.

method leads to two problems: First, it biases estimates if complete observations differ systematically from incomplete ones; second, even if an estimate is unbiased, its 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 5 shows item nonresponse rates across some selected variables.

Table 5 for example shows that – when asked about the value of their main residence – 75.5% of households provided a specific amount (third column). The other 24.5% of households are item nonrespondents, which means that either they

Table 5

Item Nonresponse for Selected Variables (Unweighted)

	Household has item		Responses by households that have the item			
	Yes	Un-known	Amount	Interval	“Don't know”/ “No answer”	Other missing values ¹
	(1)	(2)	(3)	(4)	(5)	(6)
	%					
Value of main residence ²	49.6	0.0	75.5	15.3	6.4	2.7
HMR mortgage 1: amount still owed	15.1	1.4	63.5	21.2	12.3	3.1
Monthly amount paid as rent	44.1	0.0	97.1	2.3	0.5	0.1
Other property 1: current value	12.6	0.2	74.1	15.3	9.6	1.0
Other property mortgage 1: amount still owed	1.7	0.4	70.7	7.3	14.6	7.3
Value of sight accounts	98.9	0.0	72.0	13.3	14.4	0.3
Value of saving accounts	86.0	1.6	64.6	18.6	16.0	0.8
Value of publicly traded shares	5.4	0.4	71.1	12.5	16.4	0.0
Amount owed to household	9.3	0.5	90.5	5.0	4.5	0.0
Labor status (main employment) (person 1)	100.0	0.0	99.9	0.0	0.1	0.0
Gross cash employee income (person 1)	48.8	0.1	76.7	9.9	3.4	9.9
Gross income from unemployment benefits (person 1)	6.1	0.1	83.3	9.7	6.3	0.7
Gross income from financial investments	70.9	6.6	34.3	40.7	24.0	0.9
Gift/inheritance 1: value	21.4	1.3	71.1	16.3	10.0	2.6
Amount spent on food at home	100.0	0.0	96.3	3.4	0.3	0.0

Source: HFCS Austria 2010, OeNB.

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

² Based on the HB0900 variable.

Note: HMR = household main residence.

provided a (prespecified or individual) interval (15.3%, fourth column), responded with “Don’t know” or “No answer” (6.4%, fifth column) or that their response was edited to a missing value³ (2.7%, sixth column). Nonresponse rates⁴ vary substantially across items. Variables with high nonresponse rates include e. g. questions related to outstanding balances on mortgages that use the main residence as collateral ($100\% - 63.5\% = 36.5\%$) and the household’s gross income from financial investments ($100\% - 34.3\% = 65.7\%$). With regard to the latter, 40.7% of households provided at least an interval for this type of income, which confirms the importance of asking interval questions if there is a nonresponse to a euro question. Interval questions provide valuable and often very precise information (see the online appendix and section 2.6.2 for information on the design of euro loops). Variables with low nonresponse rates include non-euro variables, such as labor status ($100\% - 99.9\% = 0.1\%$), or the amount spent on food at home ($100\% - 96.3\% = 3.7\%$).

Table 5 (second column) also shows another aspect of item nonresponse in the HFCS: There are variables, known as *branch variables* (see chart 3 in chapter 4), which may also have missing values due to nonresponses to a previous question (*head variable*) and which are thus classified as 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 in the affirmative (70.9%) are then asked the euro question; the other households, including the 6.6% of households that did not answer the yes/no question, automatically skip the euro question. As it is unknown, however, whether the 6.6% 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 taken into account as second-order (or higher-order) missing values when analyzing nonresponse to the 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 could be considerable owing to the large amount of variables with higher-order missing values. Furthermore, as complete observations usually differ systematically from incomplete ones, complete-case analysis would bias the estimates.

³ See chapter 4 for more details.

⁴ The nonresponse rate is calculated as 100% minus the value in the “amount” column in table 5.

Table 6

Logit Regression of Nonresponse in the Euro Question on Value of Sight Accounts (Unweighted)

Covariates	Coefficients
Female (person 1)	0.0775 (0.0950)
Age (person 1)	-0.0012 (0.00344)
Tertiary education level (person 1)	-0.259* (0.156)
Employed/self-employed (person 1)	-0.195* (0.113)
Residence is in Vienna	-0.194 (0.134)
Size of main residence	0.00274*** (0.000863)
Household size	0.119*** (0.0421)
Constant	-1.331*** (0.256)
Observations ¹	2,330

Source: HFCS Austria 2010, OeNB.

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

¹ The remaining 50 observations of the dataset show missing values in one of the covariates and/or filter „missing“ observations in the dependent variable and are thus not included in the regression.

For illustration purposes, table 6 shows a regression of nonresponse to the question regarding the value of sight accounts (1 if the value is missing, 0 otherwise) on several explanatory variables. We can see that item respondents differ significantly from item nonrespondents, because respondents live in smaller main residences and smaller households, have higher levels of education 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 equations⁵). It can be summarized in the following steps:⁶

- Step 1: Select the P variables Y_1, Y_2, \dots, Y_P that are going to be imputed.
- Step 2: Fill the missing values of Y_1, Y_2, \dots, Y_P with random draws from actually observed values.
- 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 right-hand side variables that is chosen from among the HFCS variables without missing values and the variables selected in step 1 (except the one being regressed); restrict the regression sample 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 values and variables with complete observations, each regression model contains a broad set of right-hand side 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 or currently imputed values (step 4). It is important that each regression model also contains a broad set of right-hand

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 (2012) 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.

side variables with *missing* values in order to preserve the joint distribution of variables with missing values. If t tends to infinity, the imputations of the 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 behind 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 are expected to converge to an approximation of M draws from the joint posterior predictive distribution of the missing values.

Although it is theoretically possible that the sequence of draws based on the above regressions 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, with data such as those from the HFCS – where there is a large number of variables, many of which have bounds, skip patterns, bracketed responses, interactions or constraints in relation to other variables – using separate regressions for each variable reflects the data better. This approach thus makes more sense than specifying a joint distribution for all variables together, as is the case under the joint modeling approach.⁷

It should be pointed out 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. This assumption is referred to as *ignorability assumption* in the literature.

Before running the above five steps, we need to prepare the data and specify all the parameters of our imputation model: e.g. the choice 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 that are going to be imputed. Our strategy is to impute as many variables with missing values as possible, which in our case amounts to around 70% of the variables with missing values. The remaining variables with missing values are not imputed with the HFCS imputation procedure because either they do not have enough variance or there are technically not enough observations to be estimated in a regression.⁸

The imputation of as many variables as possible is supposed 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 (see

⁷ See also Little and Rubin (2002) for an overview of imputation techniques.

⁸ A very small fraction (around 2%) of these variables that could not be imputed with the HFCS imputation procedure were imputed with ad hoc methods such as hotdeck after the HFCS procedure had been completed. This is because their imputation is considered very important as they are used, for example, to build important aggregate variables, such as total household income.

the description of the disadvantages of complete-case analysis in the introduction). Another important reason for adopting this strategy is that we do not want to bias the correlation structure of the data with our imputations. If we were to reject many variables for imputation, we could not subsequently use them in the regression models as right-hand side 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, a weakness 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 imputation with the variables with the least missing values and ending it with those variables that have the most missing values. Variables with the same degree of missingness are processed in a random order, but always in the same order. Head variables are always imputed before their corresponding branch variables. For example, the variable indicating whether a household has a mortgage or not is always imputed before imputing the mortgage amount, even if the degree of missingness is 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 among four different types of regression models. For quantitative variables, we use an interval regression model⁹ because all of our continuous variables are bounded either from above or from below, or from both above and below (section 5.4.6). For binary variables, we use a logit model; and for ordinal and nominal variables, we use ordered logit and multinomial logit models.

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. We have decided to use weights for the following reasons. If the regression models are misspecified, weighted regression is better than unweighted regression because it provides unbiased estimates of the regression coefficients. Weighted regression should only be forgone if regression models are correctly specified, because then it would result in inflated standard errors. As we do not know the data-generating process, we prefer to be on the safe side and use weights to arrive (with certainty) at unbiased imputations at (maybe) the cost of slightly inflated standard errors. This is especially important, as the dataset will

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.

become publicly available and we anticipate that most users will work with the imputed datasets.¹⁰

5.4.5 Variable Transformations

Before imputing variables with missing values, we transform several of them, because 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 quantitative variables involves using the logarithm. Usually, these types of variables have a highly skewed distribution; using the logarithm helps to make the distribution look closer to the normal distribution assumption which is necessary for the forecast. Another very helpful transformation for year variables is to impute durations 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 above-mentioned logarithmic transformation is carried out on the durations and not on the years.

For categorical variables, two types of transformations may be used. First, some of the nominal variables are 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 that are usually nominal are transformed into several binary variables by creating 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 quantitative 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 *loan has no set term*. When imputing such a loan-length variable it would not make sense to run the regression over these observations together with those 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 quantitative 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 quantitative and categorical variables with missing values, is that of person IDs.¹¹ Person variables are

¹⁰ Another possibility for deciding whether to use weights or not is to impute with and without weights and then compare both models. If we find significant differences between the parameter estimates of both approaches, the results suggest the use of weights to at least obtain unbiased estimates. If no significant changes between the parameter estimates are found, but instead large changes in standard errors, then this suggests an appropriate specification and minimal problems with using unweighted models. However, running the HFCS imputation procedure twice, once weighted and once unweighted, is a process which would be very time-consuming and has not been done yet. This is left for future research.

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

modeled and imputed separately for each person ID in order to avoid biased imputations (section 5.4.8); this should ensure that persons with the same IDs show relatively homogenous characteristics if they are jointly modeled. For this reason, respondents are grouped into new person ID categories, created specifically for the imputations, prior to imputation. The criteria for this categorization are the following: All male financially knowledgeable persons (FKPs), all male partners of FKPs that were second person and all other FKPs are classified as first persons (person ID = 1). All female partners of FKPs that previously already were second persons and all women that were first persons before their male partners became first persons are classified as second persons (person ID = 2). All other persons are ordered and numbered according to their age (from old to young).

In the case of households with members that engage in farming, we use a special transformation of the variables regarding the value of the household's business(es) (HD0801–HD0803) and the variable regarding the value of the household 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. In a second step, we calculate the individual variables (HD0801–HD0803 and HB0900) based on this sum and these percentages. The motivation behind this transformation is that it considerably improves the imputed values, as some households with members that engage in farming did not state the value of their main residence separately from the value of their farm but indicated the sum of both (see section 4.6.2.8 for further details).

5.4.6 Bounds

As mentioned above, we use interval regression models to impute quantitative variables in step 3 because all these variables are bounded either from above or from below, or from both above and below. 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 very unrealistic. Examples of this type of bound include nonnegativity constraints on quantitative or count variables (e.g. income or age). The lower bound for these variables is zero for all households. Furthermore, for each quantitative variable, we use the following rule: for every household, the lower bound is equivalent to half of the smallest value observed for the variable, whereas to calculate the upper bound, we take the largest observed value and double it. This helps to avoid the imputation of extreme outliers without biasing results. More examples of general bounds include share variables (e.g. share of homeownership), where we set the lower bound to zero and the upper bound to 100, or some year variables (e.g. the purchase or inheritance year of the household main residence), where the upper bound is 2011, i.e. the year in which the latest survey interviews were carried out.

Unlike general bounds, individual bounds take different values depending on each household or person; they usually ensure consistency with other variables of 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. Or looked at from the opposite angle, when imputing the total estimated consump-

tion 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 an interval (either prespecified or individual) in a euro question instead of a specific value. Such intervals are asked for after every euro question that is left unanswered; they prove very useful for imputation purposes as they yield valuable and precise information on the missing value in the euro question (see also section 5.2 in connection with table 5).

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 net excluding utilities and vice versa), loans (e.g. the initial amount of a loan is used as an upper bound for the outstanding amount of the loan 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 association between 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 non-response with respect to the variable to be imputed makes the ignorability assumption on which our imputation model relies (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, with several hundred variables, it is, however, not feasible to include all of them, as this may lead to both multicollinearity problems and computational problems. In line with Van Buuren et al. (1999) and Barceló (2006), we have 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 (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, sex and labor status of the household's first person and the latter's partner), stratification variables (e.g. province, municipality size), information provided by the interviewers (e.g. standard of living, type of neighborhood, type of building, interview atmosphere, etc.). The latter pieces of information (paradata) were extremely important for the imputations since they provided plausible explanations for item nonresponse for many variables.
2. In addition, include variables that are very good at predicting and explaining the relevant variable to be imputed. This is the classic criterion for predictors and it helps us to reduce some of the uncertainty surrounding the imputations. These predictors are identified by their correlation with the variable to be

imputed. For example, when imputing credit 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. Usually, these variables are logically connected (e.g. outstanding principal is the original loan amount minus the sum of all repaid loan amounts). 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.

3. Remove the aforementioned predictor variables that have too many missing values within 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 with percentages of observed cases within this subsample below 50% are removed and substituted 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 and age, as well as the labor and marital status of the first person.
4. Include all variables that appear in the models that will be fitted to the data after imputation. In other words, think about different economic theories that might be tested with the data and include the variables as predictors that are expected, according to these theories, to affect or explain the variable to be imputed. Failure to do so will tend to bias the results of potential users of the data 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 asset variables.

Obviously, many variables in the survey – for example, income, age or education of the first person – fulfill more than one criterion for predictor selection.

In all regression models we also include an interaction term and a main effect dummy for each one of the above predictor variables that was not inquired about in the case of every household that was asked about the variable to be imputed. For example, suppose that we want to impute a household's consumption expenditure using mortgage amount as one of our predictors. While every household in the sample was asked about consumption expenditure amounts, not all of them were asked about mortgage amounts. If, for those households that do not have a mortgage, we just set the mortgage amount to zero (corresponds to an interaction term), the estimates would be biased, because the information on whether a household has a mortgage or not would be omitted. This information should thus be additionally included as a main effect dummy in the regression model. But again, not all households were asked whether they have a mortgage, just homeowners. Thus, we should also include a homeowner dummy in the regression.

Finally, the number of predictors is restricted by the size of the subsample over 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 each one of the above four predictor categories is represented

in each regression equation (to the extent possible). Typically, the number of predictors used for each regression model is around 20% of the number of observations for the variable to be imputed for small subsamples. For large subsamples, the number of predictors usually lies between 5% and 10%. For more details on the specification of subsamples, see the next section.

5.4.8 Specification of Subsamples

Each regression in step 3 is estimated over a subsample which consists of all households and persons 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 households order mortgages by importance, which would introduce a bias in our estimates.¹²

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

5.4.9 Number of Cycles

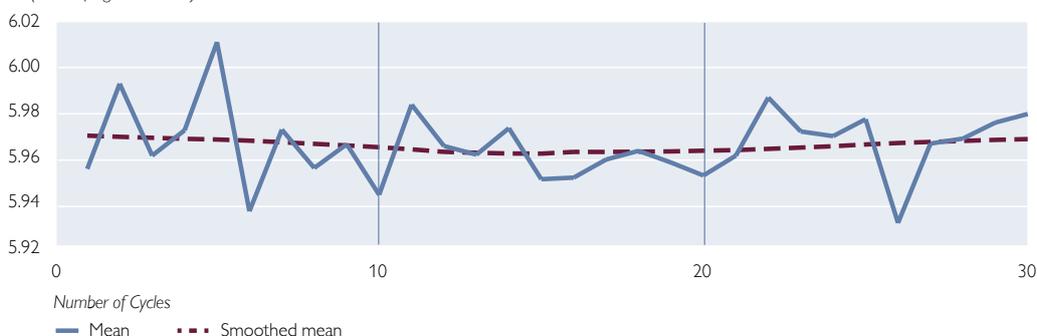
In step 4, the number of cycles (or iterations) t determines how often step 3 is repeated. As t tends to infinity, the imputed values should converge to a draw from the joint posterior predictive distribution of the variables with missing values. However, according to Van Buuren et al. (1999), in practice, convergence in these models usually occurs very quickly during the first few iterations. Given the large computational effort required for the HFCS imputation model, and following the number of iterations used in other similar surveys, e.g. the SCF (Kennickell, 1998) and the EFF (Barceló, 2006), we set the iteration number for the HFCS imputation model at $t = 6$.

Typically, we check convergence graphically by plotting the mean of the imputed values against the iteration number t ; chart 4, for example, shows this for the value of sight accounts (variable HD1110). Convergence is judged to have occurred as soon as the pattern of the imputed means becomes random. In chart 4, this seems to be the case very early on: from the first iteration on, it is no longer possible to recognize any clear trend in the smoothed curve of the imputed means of the value of sight accounts. Furthermore, chart 4 shows that the fluctuation range of the imputed means is very small, which is a further indicator of conver-

¹² Of course, in such cases, we could introduce a large number of interaction terms into our model to reduce the bias, but there still might be unobserved differences between the two groups.

Mean Value of Sight Accounts across Iterations (Unweighted)

Ln (value of sight accounts)



Source: HFCS Austria 2010, OeNB.

gence. Of course, checks of this kind can never confirm convergence (nor can any other check under the chained equations approach), but they can highlight weaknesses in the imputation model or other unusual outcomes that could be indicators of nonconvergence.

5.4.10 Number of Imputation Samples

In a 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 the 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 5. In line with the international standards set by the ECB and other similar surveys (like the SCF or EFF), we set the number of imputations at $M = 5$.

5.5 Selected Results

After imputation, the HFCS dataset is five times bigger than before, because it consists of $M = 5$ multiple imputation samples (also referred to as “implicates”). Table 7 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 the 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 88,019 before imputations. After imputations, it increases to EUR 110,526 in $m=2$, EUR 94,873 in $m=3$, EUR 190,532 in $m=4$, and EUR 125,350 in $m=5$. In $m=1$ the mean decreases slightly to EUR 87,819. Thus, on average the imputations increase the mean value of the first gift/inheritance from EUR 88,019 to EUR 121,820, i.e. by 38% (about half of the values imputed in this context are based on interval responses by households). This suggests that house-

Table 7

Means for Selected Variables before and after Multiple Imputation (Weighted)

	Mean before imputation	Multiple imputation sample means				
		<i>m</i> = 1	<i>m</i> = 2	<i>m</i> = 3	<i>m</i> = 4	<i>m</i> = 5
EUR						
Value of main residence ¹	246,203	261,468	271,337	266,286	275,096	268,015
HMR mortgage 1: amount still owed	55,745	94,427	84,670	47,135	58,619	47,657
Monthly amount paid as rent	363	334	332	335	330	334
Other property 1: current value	231,583	195,953	193,173	265,195	198,880	218,387
Other property mortgage 1: amount still owed	68,300	58,141	90,563	70,627	74,631	67,420
Value of sight accounts	2,406	3,343	3,255	3,130	2,908	3,220
Value of saving accounts	21,989	28,230	29,700	33,696	29,781	28,905
Value of publicly traded shares	30,440	23,554	36,887	22,753	28,553	22,573
Gross cash employee income (person 1)	25,871	25,075	25,254	26,517	26,230	29,403
Gross income from unemployment benefits (person 1)	6,263	6,361	6,880	6,300	6,225	6,295
Gross income from financial investments	836	800	763	730	771	787
Gift/inheritance 1: value	88,019	82,842	130,673	95,338	90,959	94,404
Amount spent on food at home	379	381	380	381	380	380

Source: HFCS Austria 2010, 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.

holds with more valuable inheritances tend to respond to this question less often than households with less valuable inheritances. The largest increases in comparison to the unimputed sample occur when imputing financial assets (e.g. the market value of stocks). Households' interval responses again play an important part here, as they provide very precise information for the imputations (see also table 5).

However, for some variables, the mean does not change significantly; 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 (table 5). 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 7 also shows that the uncertainty of imputations can vary a lot depending on the variables. For some variables (e.g. mortgage 1 that uses the main residence as collateral), 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 means 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 seen 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 procedure known as *multiple imputation by 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 of the HFCS in Austria was specifically developed by the OeNB in collaboration with the Institut für empirische Sozialforschung GmbH – IFES. Sampling means selecting 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) will 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 assured 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 is not avoidable, sampling – together with imputation and weighting – serves to produce the best unbiased estimates (and confidence intervals) and to keep uncertainty as low as possible under restrictions like costs, time and practicability. Therefore, a sound and clear sampling design is essential and the starting point for any survey.

This chapter 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 sample population (section 6.5), which build the main part of the sampling procedure. Section 6.6 contains concluding remarks.

6.2 Target Population and Sample Frame

The first step in determining the sampling procedures 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 as separate households¹.”

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
5. Persons absent for long periods, but having household ties: persons working away from home
6. Persons temporarily absent but having household ties: persons in hospital, nursing home, boarding school or other institution”

¹ See ECB (2011, p. 7).

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

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

Additionally the Austrian HFCS does not cover homeless people. People without a residence could not be reached with the survey as sampling was based on dwellings (see below). 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 the complete list of private mailboxes in Austria as our frame population. 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 the sample – i.e. households for our purposes – are drawn from all the desired parts of the population. Stratification in the Austrian HFCS was done by geography (based on NUTS-3 regions¹) and by municipality size categories.
- “Two-stage cluster sampling” means that, in a first step, primary sampling units (PSUs) are selected from each geographical unit (i.e. stratum) and that, in a second step, secondary sampling units (SSUs) are drawn within each PSU. The two-stage sampling design of the Austrian HFCS (see the chart in this box) consists, first, of selecting a random sample of enumeration districts (the smallest geographical unit for which statistical data are available) from each stratum and, second, of selecting a random sample of households (mailbox

Overview of Two-Stage Sampling



Source: Statistics Austria, HFCS Austria 2010, OeNB.

¹ See www.statistik.gv.at/web_de/klassifikationen/regionale_gliederungen/nuts_einheiten/index.html (accessed on January 22, 2013). 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.

codes) from each sampled enumeration district. The households thus constitute the secondary sampling units (SSUs). The two-stage cluster design reduces costs due to relatively small distances between the 12 households (8 households in Vienna) selected within each PSU while it ensures a sufficiently high number of PSUs within the individual stratum.

Thus we ensure that households from every single stratum are invited to take part in the survey. In total, the gross sample of the Austrian HFCS consists of 170 strata, 422 PSUs and 4,436 households.

6.3 Background – The (External) Data Sources

Given the definition of the target population, proper information on the geography and on households in Austria is needed. A perfectly representative draw of our sample presupposes that the target population is represented correctly by the frame population. 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). This theoretical optimum cannot be achieved in practice. The HFCS in Austria has been designed to meet this goal subject to the constraints of the available data sources. The following sections provide more details on the data on which sampling was based 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 data from the Post Office for the actual draw of SSUs (Secondary Sampling Units), namely households. 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 patterns of Austria, i.e. data on the NUTS-3 regions, and the enumeration districts (PSUs) from the population census 2001.² These enumeration districts are the smallest territorial units in Austria for which basic data characteristics are collected by default by Statistics Austria (each enumeration district contains around 410 dwellings on average).

In addition, we relied on the municipality directory of 2006 (to categorize the municipalities) and on the Micro census of 2009 (for the population data).³ Both datasets are collected and provided by Statistics Austria. The Micro census 2009 indicates in particular the population size 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.

² See www.statistik.at/web_de/statistiken/bevoelkerung/volkszaehlungen_registerzaehlungen/index.html (accessed on January 22, 2013).

³ See *Statistics Austria (2011a)*.

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 on mailbox codes 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 mailboxes in the respective estate. Specifically, we used a commercial product called *Adress.Certified* developed by the Austrian Post Office. This address register captures information at the building level (including the information on street, house number and whether the estate is privately or commercially used). It can be bought in combination with a product called *DATA.DOOR*, which is a directory of mail delivery points (abbreviated PAC, for post office-certified address codes, in table 10), i.e. a directory of all addresses (mailbox codes) in Austria to which mail can be delivered. This information is available at a disaggregated level. In Austria, there are about 3,930,000 private mail delivery points and 233,000 commercial mail delivery points.

Thus, our starting point was some 3.9 million private mail addresses. Very few remaining commercial addresses had to be withdrawn after the first contact by the interviewer (e.g. if the interviewer arrived at the address and noted that this was an invalid address or a commercial building; see chapter 4) and were given weights of zero since they do not belong to the target population.⁴

The post office data we used do not reflect whether a given address is a household's main residence or not. Yet they provide a realistic picture of households and thus meet the HFCS requirement of reflecting actual living situations. In particular, unlike other data sources the post office data cover households at addresses that are registered only as a secondary home but fulfill the HFCS definition of a household. They have thus been included in the frame population because they have a post-certified address code.⁵ At the same time, secondary homes inhabited by households whose main residence address was identifiable as such were removed in order to ensure that every household in Austria appeared only once in the remaining list of the post-certified address codes (see also chapter 7).

6.3.3 Profile.Address and IFES

To identify the names of the households that correspond to the selected mailbox codes – information that is not evident from the datasets described so far – the polling firm, IFES, relied on its databases or bought the corresponding addresses from a commercial provider called “Profile.Address.”

This information was needed in the contact phase when households received personal letters of invitation to participate in the survey.⁶

⁴ Having deducted these commercial addresses, the total sum of weights is around 3.77 million; hence the total estimated number of households in Austria is 3.8 million.

⁵ The post-certified addresses for some 3.9 million households compare with about 3.6 million household addresses documented by other sources (such as the Micro census based on the residence registry).

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

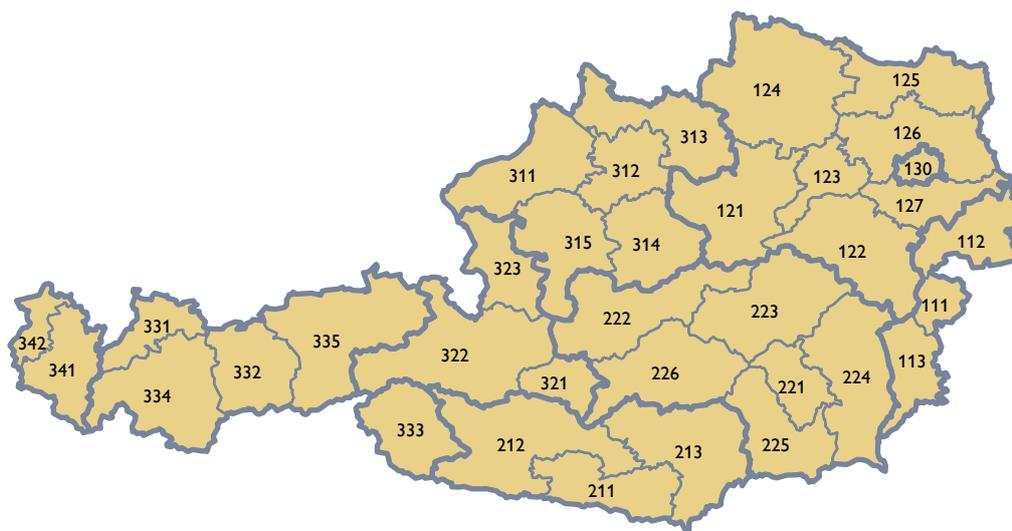
6.4 Stratification and Sample Size

6.4.1 Stratification

The Austrian HFCS essentially used two indicators for stratification, the first one being the 35 NUTS-3 regions (chart 5).

Chart 5

NUTS-3 Regions



Source: Statistics Austria.

With the exception of the capital city Vienna, each NUTS-3 region was divided further by the following 8 municipality population size categories.

- 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 the urban areas with the provincial capitals. Vienna stands out as it is the national capital and the only city with more than 1 million inhabitants; it was simply subdivided into its 23 districts.

This very fine stratification yielded 193 strata. Where the number or rather proportional share of households per stratum was too small to have allowed the selection of enumeration districts, individual strata were merged with neighboring strata to increase the share of households and hence insure the selection of PSUs within each stratum. This exercise left the HFCS with 170 strata that covered all households in Austria. See table 8 for the distribution of strata across provinces and municipality size categories.

Table 8

Allocation of Strata within the Sample

Province	Municipality size (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 Mio	over 1 Mio	
Vienna	0	0	0	0	0	0	0	23	23
Lower Austria	7	7	7	7	5	5	1	0	39
Burgenland	2	2	2	1	1	0	0	0	8
Styria	5	5	5	5	3	1	1	0	25
Carinthia	3	3	1	3	2	1	1	0	14
Upper Austria	5	5	4	5	3	2	1	0	25
Salzburg	2	2	2	2	2	0	1	0	11
Tyrol	3	3	3	4	3	0	1	0	17
Vorarlberg	1	1	2	1	2	1	0	0	8
Total	28	28	26	28	21	10	6	23	170

Source: Statistics Austria (municipality directory 2006).

Each stratum contained about 50 PSUs on average, which in turn contained around 470 households on average (see also section 6.5.1).

6.4.2 Sample Size

As a rule, the variance of estimates based on the underlying data will be the smaller, the larger the sample is. At the same time, the cost of data collection increases with size. This trade-off has to be balanced to yield reasonably precise estimates subject to the given budget constraints. Furthermore, given the focus of the survey and the analyses HFCS was supposed to facilitate, 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 frame population) and to provide some insights into the regional differences within Austria. We know from previous OeNB surveys (Household Survey on Housing Wealth, 2008, and Survey of Household Financial Wealth, 2004) that at least 2,000 households need to be successfully interviewed and that the unit nonresponse rate can be expected to reach some 40% (with expected differences between Vienna and the rest of Austria).⁷ With some room for extraordinary events, the HFCS survey was therefore designed to yield a sample of successful interviews of 2,700 households and a participation rate of 50% for Vienna and 65% for the rest of Austria. These participation rates are rough estimates based on the experience of past surveys.

The net sample of $n = 2,700$ was divided between the nine provinces, based on their population shares observed by Statistics Austria in the Micro census of 2009⁸ (table 9). These figures, which corresponded to the targeted number of secondary sampling units (SSUs), were translated in a next step into gross samples of SSUs based on the estimated participation rates. Within each enumeration district, 8 households were subsequently drawn in Vienna and 12 households in the rest of

⁷ The participation rate per stratum in the first HFCS wave will have to be taken into account in future surveys.

⁸ Although the Micro census collected by Statistics Austria in 2009 is based on a different definition of households than the HFCS, we used this information as an approximation.

Table 9

Identification of the Number of PSUs to Be Drawn

Province	% of households	Target sample	Target participation rate in %	Gross sample ¹	Number of households per PSU (enumeration district)	Number of PSUs to be drawn
Vienna	23	629	50	1,258	8	157
Lower Austria	18	499	65	767	12	64
Burgenland	3	84	65	129	12	11
Styria	14	376	65	578	12	48
Carinthia	7	179	65	275	12	23
Upper Austria	16	437	65	672	12	56
Salzburg	6	168	65	258	12	21
Tyrol	8	216	65	333	12	28
Vorarlberg	4	113	65	174	12	14
Total	100	2,700				422

Source: Statistics Austria (municipality directory 2006), HFCS Austria 2010, OeNB.

¹ In sum, the targeted gross sample does not match 4,436 exactly due to rounding.

Austria, with the respective difference reflecting the shorter distances between buildings in Vienna. Thus, we arrived at the PSUs to be drawn in each province.

In sum, the Austrian HFCS sample design produced 422 PSUs across all strata and a gross sample size of 4,436 households that were invited to participate in the HFCS (see box 2 in chapter 7 for information on the number of successfully interviewed households). Drawing possible substitute addresses was explicitly excluded in the HFCS to begin with, in order to ensure that all households of the gross sample would be interviewed with the same commitment so as to prevent data distortions (see also section 4.4.1).

6.5 The Two Stages of the Random Draw

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

- stage one: random draw of PSUs (enumeration districts) from each stratum
- stage two: random draw of a predefined number of households (mailbox codes) from each PSU

6.5.1 First Stage

As the PSUs of the Austrian HFCS we chose the smallest territorial units in Austria, the so-called enumeration districts, of which there are 8,745 in Austria. On average, an enumeration district contains 410 households, but there are also PSUs with only very few households. These 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 8,407, which then had 470 households on average. Each PSU has the same probability to be drawn (irrespective of size) from a given stratum, i.e. there is no “probability proportional to size” selection – a PSU comprising numerous households has the same probability to be chosen within a stratum as a PSU with fewer households.

From the description above it is clear that the number of PSUs to be drawn in each province is determined a priori by the chosen sample size. To translate the numbers allocated to each province (table 9) into the desired number of PSUs within a given stratum, the former were divided up in proportion to the number of households in that stratum. For example, the 56 PSUs to be drawn in Upper Austria (table 9) were divided up among the 25 strata in this province according to population shares.

With the number of PSUs to be drawn per stratum having been established, they were selected in next step by simple random sampling without replacement.

6.5.2 Second Stage

With 422 PSUs having been randomly drawn, we turn to the second stage in which households are selected.

As outlined above, eight or twelve mail delivery points were randomly selected from each PSU drawn, depending on whether the PSU was domiciled in Vienna or elsewhere in Austria. In this process, every household in a given PSU had an equal probability of being selected into the sample. This procedure resulted in a gross sample of 4,436 households in Austria.

6.5.3 Practical Implementation

Table 10 illustrates how the second stage (following the draw of PSUs) worked: Austrian Post Office data were used to arrive at the appropriate mail delivery point, which essentially contains the address but not the holder of this address. To identify the household's name corresponding to the selected mailbox code the survey firm, IFES, relied on its databases or, where necessary, bought the corresponding name from a commercial provider called "ProfileAddress."

Table 10

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 a personal letter signed by the governor of the OeNB. This letter contained information on the survey and an invitation to take part (chapter 3).⁹

⁹ See the online appendix for the invitation letter.

6.6 Concluding Remarks

This chapter provides information on the sampling design specifically developed for the HFCS in Austria that has now been applied for the first time. As described, the survey is based on stratified clustered two-stage cluster random sampling, consisting of a random draw of primary sampling units (here, enumeration districts) from each stratum plus a random draw of households (mailbox codes 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 two aspects being particularly important:

- Sampling is based on a broad-based concept of the target population (without differentiating between main residences and second homes), as a result of which all households that correspond to the HFCS household definition have a positive probability of being selected.
- The stratification structure is very fine, as a result of which all segments of the Austrian population are represented in the survey.

At the same time, since one of the topics covered by the HFCS is the wealth structure of households one would want to oversample certain groups of the population such as indebted or wealthy households in order to improve the efficiency of estimates of these subgroups. However, the underlying information needed for such oversampling is unfortunately not available as yet.

7 Construction of Survey Weights

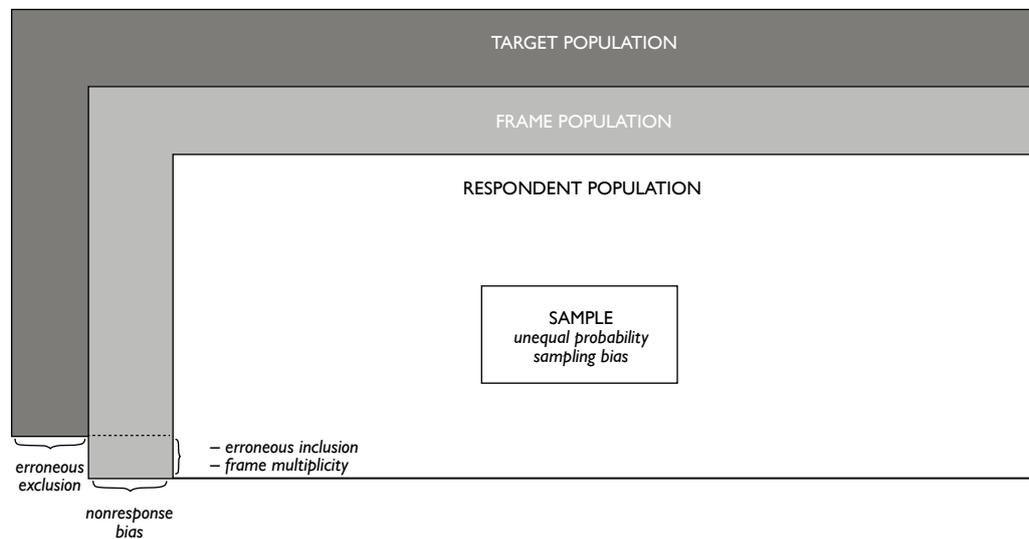
7.1 Introduction

Survey weights are usually constructed for two reasons: first, to make the sample representative of the target population and second, to reduce sampling variance.

The target population of the HFCS consists of all households in Austria, with a household being defined as a person 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: the unequal probability sampling bias, the frame bias (due to an imperfect frame population), and the nonresponse bias (chart 6).

Chart 6

Misrepresentation of the Target Population in the Sample



Source: Adapted from Biemer and Christ, 2008

The unequal probability sampling bias is due to the fact that not every household has the same probability of being selected into the sample. For example, given the sampling design of the HFCS, households in small primary sampling units (PSUs) have a higher probability of being selected into the sample than households in large PSUs. Another example is the oversampling of Viennese households in the HFCS. To correct these misrepresentations, we construct design weights, which will be explained below. For further details about the sampling design of the HFCS see chapter 6.

The frame bias is caused by imperfections in the survey frame from which the sample is drawn. In the HFCS, the frame population is a list of all personal postal addresses in Austria (chapter 6). Erroneous exclusion of households could imply an imperfection with respect to the target population. In other words, we may have excluded, households without a postal address, for example, one-person households living together in residential communities and sharing an address that contains only one of these households. Then these types of households would be

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.

underrepresented. Another imperfection of the frame could be erroneous inclusion, that is, including addresses not belonging to households, for example, those of companies² or households 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 constructing design weights³ (in case of erroneous inclusion and frame multiplicity) or by using poststratification weights (in case of erroneous exclusion). In this chapter, we explain the construction of these weights.

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, which is commonly found in the literature (see e.g. Kennickell and McManus, 1993). Thus, estimates for the entire frame population would be biased with respect to these group characteristics even though they are unbiased for the responding population. Using nonresponse weights can correct this bias (section 7.2.3).

Furthermore, as mentioned at the beginning, survey weights can help to reduce sampling variance, and, hence, increase the precision of the estimators. Ideally, the precision of the estimators should be improved via 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 our estimators, have not been 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 realized had these variables been used for stratification can be realized by poststratifying by these variables when constructing the poststratification weights. These poststratification weights were also employed for correcting erroneous exclusion (7.2.4).⁴

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

7.2 Construction of Survey Weights

7.2.1 Weight Components

We aim to construct a final survey weight w_i for each 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, there are various reasons why households may misrepresent the target population. Depending on the type of misrepresentation, a specific adjustment by applying weights is required. In the HFCS, three types of weights are used: design weights w_{Di} , nonresponse weights

² Although the frame population was cleaned of addresses of companies, there may still be some erroneously left.

³ Sometimes referred to as noncoverage weights.

⁴ Poststratification weights can also correct for a third type of sample-specific bias: the target population may be randomly misrepresented by the specific households drawn into the sample.

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 person level instead of household level, we do not construct weights for persons because the main focus of the survey is on the household level.⁵

7.2.2 Design Weights

Design weights help to reduce the unequal probability sampling bias, to repair erroneous inclusion and to correct for frame multiplicities. In the HFCS, we construct the design weights in three steps:

- Step 1: Constructing a preliminary set of design weights to eliminate the unequal probability sampling bias
- Step 2: Calibrating this preliminary set of design weights in such a way that they sum up to the frame population total
- Step 3: Resetting the calibrated design weights to zero for those observations that were either erroneously included or that were duplicated

In order to eliminate the unequal probability sampling bias, in step 1 we had to put proportionally more weight on households with a lower probability of being selected into the sample than other households and vice versa. Therefore, we had to calculate the probability of selection for each household and then get the preliminary design weight by inverting this probability. For example, a household with a probability of selection equal to 0.001 got a preliminary design weight equal to $1,000 = 1/0.001$, which is much higher than that used for a household with a probability of selection equal to 0.009, which would be $111 = 1/0.009$.

When calculating households' selection probabilities, we had to keep in mind the two-stage sample design: in the first stage the smallest territorial units, enumeration districts (Zählsprenkel), were drawn, and in the second stage households within the enumeration district were drawn (chapter 6). The probability that the i^{th} household in the j^{th} enumeration district is selected into the sample is the product of the enumeration district selection probability and the probability that the household is selected given that the household's enumeration district is selected. The inverse of this product is the preliminary design weight.

Due to this household reweighting, the estimated frame population total would no longer correspond to the actual frame population total. Therefore in step 2, the preliminary set of design weights was calibrated according to the actual frame population total.

Finally, although the frame population had been carefully prepared and cleaned before sampling, there were still some ineligible (see box 2) or duplicated observations in our sample (see also section 4.6.2.12), for example addresses belonging to companies or to care residences, or multiple domiciles of commuters. We flagged all such cases detected during the field work as ineligible or duplicated in our sample by setting the calibrated design weights equal to zero.

⁵ In order to carry out analyses at the individual level, additional weights would have to be constructed. This is planned as a future research project in the HFCS.

Table 11 shows some statistics of the obtained HFCS design weights across Austria's federal provinces. Vienna is the federal province with the lowest median weight, which makes perfect sense as households living in this province are oversampled and would have created a bias had they not been re-weighted downward by the design weights.

The value of a household's design weight can be interpreted as the number of households in the frame population that is represented by this household. For example, the median Viennese household represents 693 households in the frame population. Thus, the size of the frame population can be estimated by summing up all design weights, which in our case is 3,773,956 households.

Table 11

HFCS Design Weights by Federal Province

	Mean	Median	Minimum	Maximum
Vienna	747	693	0	2,554
Lower Austria	931	882	0	2,437
Burgenland	916	1,050	0	1,379
Styria	902	756	0	3,271
Carinthia	874	864	0	2,229
Upper Austria	844	857	0	2,394
Salzburg	905	969	0	1,712
Tyrol	872	879	0	1,852
Vorarlberg	895	849	0	1,560
All	851	833	0	3,271

Source: HFCS Austria 2010, OeNB.

Box 2

Unit Nonresponse in the HFCS in Austria

In the HFCS in Austria, successful interviews were conducted with 2,380 households of the gross sample, which comprised 4,436 addresses. The remaining 2,056 addresses were either classified as unit nonresponse (1,888 households), ineligible addresses (162 addresses) or addresses of unknown eligibility (6 addresses).

The unit nonresponse cases are households as defined in the HFCS that were not interviewed successfully. There were several reasons for this. First, and most importantly, households actively refused to take part in the survey, either by refusing to be interviewed or by breaking off the interview or by failing to keep the appointment for the interview and refusing to be contacted again. This applied to a total of 1,693 households. Second, no contact at all could be established with 75 households. Third, there were other reasons for nonresponse, such as illness, language barriers or ex post exclusion of interviews due to a high number of missing or unreliable values. This applied to 120 households.

In addition, 162 addresses were classified as ineligible because they were not part of the target population, for instance addresses of companies, empty buildings or second homes of households that were reached via their main residence address. Furthermore, there were another six addresses whose eligibility status was impossible to ascertain as the interviewers were unable to reach or find them. In accordance with the distribution of the eligibility status of the rest of the observed addresses in the sample, one of the six addresses was randomly chosen to be ineligible and the remaining five to be eligible.

The eligibility rate in the HFCS sample finally came to 96.3% and the nonresponse rate of the eligible households amounted to 44.3%. This means that successful interviews were conducted with 55.7% of the eligible households in the HFCS sample. 39.6% 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 frame population and the sample is nonresponse biased (chart 6). In the HFCS, this is indeed the case, as can be seen in table 12.

Table 12

Response Propensity Estimates Based on a Logit Regression Model

Covariates	Coefficients
Dwelling rating low	-0.260 (0.233)
Dwelling location countryside	-0.077 (0.0837)
Dwelling location Vienna	-1.091*** (0.0924)
Female interviewer	-0.138* (0.0738)
Interviewer's age	-0.005 (0.00359)
Interviewer's second-stage tertiary education	-0.550*** (0.0906)
Interviewer's working experience (in months)	0.00172*** (0.000512)
Number of interviewer contacts	-0.022 (0.0376)
Constant	0.883*** (0.197)
Observations ¹	4,273

Source: HFCS Austria 2010, OeNB.

¹ The remaining 163 observations in the data set are ineligible and are therefore not included in the regression.

Note: Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

The table shows a logit regression of household participation in the survey (1 if the household participated, 0 otherwise) on a set of variables that explain participation in the survey. The results show on the one hand that households living in Vienna and households contacted by interviewers with a high level of education have a lower probability of participating. On the other hand, households contacted by experienced interviewers achieve a significantly higher probability of responding than households contacted by other interviewers. This suggests that nonresponse is not random.

This bias can be corrected by using nonresponse weights, i.e. by attaching a higher nonresponse weight to households with a low probability of responding than the weight attached to households with a high probability of responding. To calculate the response propensities and the corresponding nonresponse weights we combine the weight-

ing class adjustment method with the model-based adjustment method (Biemer and Christ, 2008). The algorithm can be summarized in the following three steps:

- Step 1: Using the logit regression model shown in table 12 to estimate the probability of response for each household (given that the household was selected into the sample)
- Step 2: Grouping these households' response propensities into 5 classes (quintiles) and calculating the average response propensity for each class (unweighted total number of respondent households/unweighted total number of households)⁶
- Step 3: Getting the nonresponse weight of each class by inverting the average response propensity of each class

The advantage of this approach is that this method stabilizes the nonresponse weights because the response propensities predicted by the regression model vary widely and may contain extreme values.⁷ Information collected through interviewer questionnaires (e.g. level of education and experience of the interviewer)

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

⁷ Another problem of the simple use of logistic regression modelling, pointed out by Iannacchione et al. (1991), is that it does not ensure that the weighted sample marginal distributions conform to the population marginal distribution.

Table 13

HFCS Nonresponse Weights by Response Propensity

Quintile	Predicted response propensity %	Weight
First quintile	0 to 40	2.837
Second quintile	40 to 61	2.110
Third quintile	61 to 64	1.580
Fourth quintile	64 to 67	1.471
Fifth quintile	67 to 100	1.457

Source: HFCS Austria 2010, OeNB.

was found to correlate strongly and statistically significantly with households' response propensity and was therefore used in step 1.

The HFCS nonresponse weights are shown in table 13. They took five different values, one for each response class, and, by construction, those households with a high response propensity got a lower weight than those with a low response propensity. Nonrespondent households got a nonresponse weight equal to zero.

7.2.4 Poststratification Weights

Erroneous exclusion may be – as mentioned in the introduction – an imperfection in the HFCS frame with respect to our target population. We may have missed households without postal addresses, which implies that these types of households would be underrepresented. If there existed an external data set that covered these households and all the other ones of our target population, we could use it to adjust our sample to this external data set; we could put more weight on households that are similar to those without postal addresses in such a way that the estimated size of the target population in the HFCS would be the same as the one in the external data set.

Unfortunately, such a data set does not exist in Austria. Similar ones, like the EU-SILC (EU Statistics on Income, Social Inclusion and Living Conditions) or the Austrian Micro census, target a different population of households, which is 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 Micro census include only households living in a dwelling officially registered in the Austrian residence registry (Zentrales Melderegister) as their main residence. This definition excludes a subset of households which are 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 why households' actual main residences are not registered as such. For instance, students that studying away from home may keep their main residence at their parents' address even though they are already a household on their own by HFCS definition; others may just forget to register the address at which they actually live as their main residence. These and other problems when using main residence addresses for sampling households via the Austrian residence registry are also acknowledged by Statistics Austria.⁸

Given that these data sets also suffer from erroneous exclusion, it does not make sense to reweight the sample according to the target population *size* of these data sets. In addition, the erroneous exclusion bias in the HFCS sample is likely to be very small, as the vast majority of households do have postal addresses. How-

⁸ For the Micro census see Statistics Austria (2011a), page 10; for the EU-SILC see Statistics Austria (2011b), page 27.

ever, it appears to make sense to reweight the HFCS sample according to the *structure* of the target population of the external data sets. In particular, this may deliver a better picture with regard to the proportions of households in the Austrian provinces as the Austrian Micro census uses a much larger sample than the HFCS. We constructed poststratification weights that put more weight on households with a low probability of being included in the frame and less weight on households with a high probability; we did not change the HFCS frame population size, but we adjusted the proportions of various household types. This increases the comparability between the HFCS and the Micro census 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 (section 7.1).

The HFCS poststratification weights are constructed following the poststratification cell adjustment method (Biemer and Christ, 2008) and using the Austrian Micro census data available during the field phase of the HFCS in Austria (2010 Q4). The procedure was as follows:

- Step 1: Choosing good predictors for a household being included in the HFCS frame and cross-tabulate these variables to build the poststratification cells; our poststratification variables are municipality size and household size.
- Step 2: Calculating the average inclusion-in-frame propensity for each cell

$$\text{Average inclusion-in-frame propensity} = \frac{\text{HFCS frame population in the cell}}{\text{Total HFCS frame population}} \cdot \frac{\text{Micro census frame population in the cell}}{\text{Total Micro census frame population}}$$

- Step 3: Getting the poststratification weight by inverting the average inclusion-in-frame propensity for each cell.

We grouped household size into two very broad classes: 1 to 4 persons, and 5 or more persons. This way larger households are not underrepresented in the HFCS sample and, at the same time, we did not create a bias in our target population toward the Micro census' target population, because we are aware of the fact that in the HFCS target population the proportion of smaller households is higher (as it

includes households living in a dwelling without official main residence status in the residence registry). Moreover, we assigned the households to seven municipality sizes.

The HFCS poststratification weights are shown in table 14. They take 14 different values, one for each municipality size class and household size class. We can see that large households were underrepresented in the HFCS frame population because they tend to have higher poststratification weights. Table 14 also shows that large municipalities have lower poststratification weights and were therefore overrepre-

Table 14

HFCS Post-Stratification Weights for Municipality Size Class and Household Size Class

Municipality size (number of inhabitants)	Household size (number of persons)	
	1 to 4	5 or more
Up to 2,000	0.974	1.133
2,001 to 3,000	1.169	1.055
3,001 to 5,000	0.963	1.447
5,001 to 10,000	1.053	1.813
10,001 to 20,000	0.915	2.012
20,001 to 1 Mio	0.911	1.913
Over 1 million	0.977	1.249

Source: HFCS Austria 2010, OeNB.

sented (the average weight of the last three size classes is lower than that of the first four ones).

7.2.5 Final Weights

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

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

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

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

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.

Table 15 summarizes the HFCS nonresponse and poststratification adjustments, which we obtain by taking the product of the nonresponse weights (table 13) and the poststratification weights (table 14). These adjustments take 48 different values depending on household size, municipality size and predicted response propensity classes.

Table 15

Nonresponse and Poststratification Adjustments

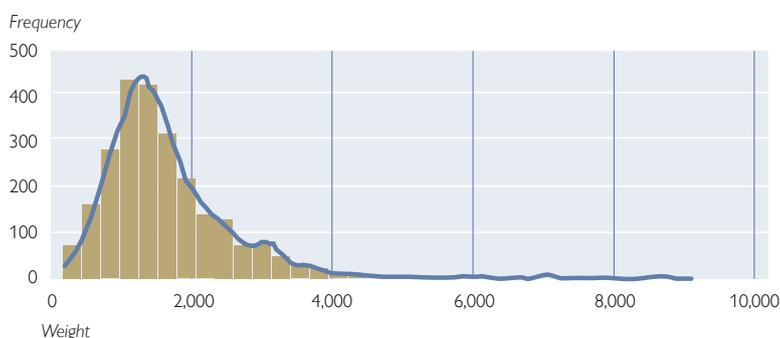
Poststratification cells		Nonresponse cells				
		0 to 40	40 to 61	61 to 64	64 to 67	67 to 100
Up to 2,000	1 to 4		2.055	1.539	1.432	1.419
	5 or more		2.392	1.791	1.667	1.651
2,001 to 3,000	1 to 4		2.468	1.848	1.720	1.703
	5 or more		2.227	1.668	1.552	1.537
3,001 to 5,000	1 to 4		2.033	1.523	1.417	1.403
	5 or more		3.053	2.286	2.128	2.107
5,001 to 10,000	1 to 4		2.222	1.665	1.549	1.534
	5 or more		3.826	2.866		2.641
10,001 to 20,000	1 to 4		1.931	1.446	1.346	1.333
	5 or more			3.181	2.960	2.931
20,001 to 1 million	1 to 4		1.923	1.440	1.340	1.327
	5 or more			3.024	2.814	2.787
Over 1 million	1 to 4	2.771	2.061			
	5 or more	3.544				

Source: HFCS Austria 2010, OeNB.

Note: Empty cells denote that there are no households of this size in the sample.

Chart 7

Distribution of Final HFCS Weights



Source: HFCS Austria 2010, OeNB.

Table 16

Comparison of Weighted and Unweighted Means of Selected HFCS Variables (Imputed)

	Mean	
	unweighted	weighted
Household size (number of persons)	2.11	2.13
	<i>Percentage share of all households</i>	
Vienna	28.3	23.3
Lower Austria	17.3	16.0
Burgenland	3.0	4.1
Styria	13.0	15.4
Carinthia	6.2	6.8
Upper Austria	15.1	15.3
Salzburg	5.7	6.6
Tyrol	7.6	8.4
Vorarlberg	3.8	4.1
	<i>EUR</i>	
Estimated household monthly net income	2,287	2,306
Value of main residence ¹	257,666	258,072

Source: HFCS Austria 2010, OeNB.

¹ The variable DA1110 was used here.

Finally, when taking into account also the design weights, we obtain the HFCS final weights, whose distribution is shown in chart 7. The HFCS final weights range from 169 to 9,054, with the mean being 1,586 and the median 1,429. Their distribution is slightly skewed to the right, which is not atypical for unequal probability sample designs: Households with a higher probability of selection (design weights below average) dominate the sample. This effect is reinforced by the further weight adjustments.

7.3 Selected Results

Table 16 shows the impact that the HFCS final weights have on estimations by comparing selected weighted and unweighted mean values of HFCS variables. For example, we can see that Viennese households were strongly down-weighted from 28.3% to 23.3%. This means that despite their high non-response rate, in sum, Viennese households were clearly overrepresented in the sample with respect to the target population. We can also see from the table that households with more expensive primary residences were underrepresented in the unweighted sample, which can probably be explained by these households' higher nonresponse rates.

The use of the final HFCS weights is sufficient when calculating e.g the weighted statistics shown in table 16.

To calculate the corresponding correct

variances or standard errors of these estimators, however, resampling weights, which are described in chapter 8, are necessary.

7.4 Concluding Remarks

We constructed a set of final weights to correct imperfections in the unweighted HFCS sample with respect to the HFCS target population. These imperfections are the 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. 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 27.7% ($UWE=1+cv^2=1.2768$) as a result of weighting, which is a relatively small effect. Therefore, it is not necessary to apply weight trimming methods. Furthermore, we prefer to have much less bias at the cost of a small increase in variance in order to avoid distorted results being wrongly considered significant too often.

An explanation of how to correctly use the weights in Stata[®] is provided in chapter 9, User Guide

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. To calculate the corresponding correct variances or standard errors of these estimators, however, replicate weights, which are described here, are necessary.

HFCS sampling involves a variety of features such as stratification or cluster-sampling. 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 (Kolenikov, 2010).

Unfortunately, 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 weights methods* (also called *replication* or *resampling methods*) and *linearization*.¹

Until recently, the literature preferred linearization to replication, mostly because it requires less computational power. However, an important disadvantage of linearization is that confidentiality provisions do not allow the inclusion of all information necessary for linearization. One way to avoid the problem that certain information is not available for reasons of privacy 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 data set – e.g. stratum and primary sampling unit (PSU) variables – it is virtually impossible for the data user to identify a given respondent (Stata Library, 2012).

Moreover, the linearization method cannot be used for estimating the variance of nonlinear statistics (medians, quartiles, etc.) as it requires computing derivatives of smooth functions; however, quantile functions, for instance, are discontinuous. Replicate weights, by contrast, are well suited for estimating the variance of such statistics (Heeringa et al., 2010).

Given the above-mentioned confidentiality constraints and because the HFCS data in particular facilitate the analysis of distributional parameters such as medians and quartiles, we decided that the variance estimation method to be employed for the HFCS should be based on the use of replicate weights. In the following section we describe how replicate weights are constructed for the HFCS in Austria.

8.2 Construction of Replicate Weights

8.2.1 The Replication Method

The replication method aims at estimating the variance of an estimated population parameter. The idea behind it is to estimate, in a first step, population parameters for individual subsets (so-called replicates) of the sample observations. In a second

¹ For a comprehensive overview of variance estimation methods, see Levy and Lemeshow (2008) or Heeringa et al. (2010).

step, the variability of these estimated population parameters across all replicates is calculated, resulting in the desired variance of the overall estimated population parameter (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 (Kolenikov, 2010). The replicate weights $w_i^{(r)}$ for $r=1, \dots, R$ will be published together with the HFCS data set.

There are different methods to form such replicates. The three major replication methods used in the survey literature are *balanced repeated replication*, *jackknife* and *bootstrap replication*. Although in most cases, the estimators of all replication methods converge to one another as the sample increases, simulation studies have shown that for quantile estimation bootstrap and balanced repeated replication are better suited than the jackknife (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 extended 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 to the sampling distributions of the relevant statistics.

8.2.2 Sampling Error Calculation Model

To mimic the original sampling procedure, we build a sampling error calculation model which is a simplification of the actual complex sample design (chapter 6). Such models are used because, in many cases, the complex sample design intended for data collection poses analytical problems for variance estimation (Heeringa et al., 2010).

In the HFCS in Austria, one necessary simplification of the sampling error calculation model compared with the original sampling procedure is to collapse 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: In 81 strata out of a total of 170, only one PSU was drawn. For the sampling error calculation model, every single-PSU stratum is paired with the geographically nearest stratum to form a single pseudo-stratum. Although collapsing strata produces an upward bias in the estimated variance, this bias is kept as small as possible by collapsing geographically close strata and thus keeping the PSUs within one pseudo stratum very homogeneous. In this context it must be pointed out that upward biases of standard errors lead to a loss in statistical power. In general, however, this is more acceptable than negative biases of standard errors, which lead to results that will too often qualify as statistically significant.

Table 17 shows how stratum size (number of PSUs drawn per stratum) changes when the HFCS sampling error calculation model is used instead of the original

HFCS sample design: When collapsing strata in the sampling error calculation model, their number decreases from 170 to 89, which means stratification is still very high. Moreover, the mean stratum size increases from 2.5 PSUs to 4.7 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 a drawn PSUs are selected into the replicate sample.

The sampling error calculation model does not require any further simplifications. In particular, the nonresponse and poststratification weight adjustments are implemented exactly as in the original weighting procedures (chapter 7), and a finite population correction² is performed.

8.2.3 Construction of Replicate Weights

The algorithm used to construct the HFCS replicate weights comprises the following steps:

- Step 1: Draw m_h PSUs with replacement within each pseudo-stratum h .
- Step 2: Adjust the final survey weights of the drawn observations to create a new set of replicate weights. In particular, apply the same nonresponse and poststratification weight adjustments (sections 7.2.3 and 7.2.4) as for the 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 is a common choice that ensures the efficiency of the bootstrap estimators without leading to violations of the natural parameter ranges (Kolenikov, 2010).

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

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

Table 17

Comparison of HFCS Design Strata and HFCS Pseudo Strata

	Design strata	Pseudo strata
Number of strata	170	89
Mean size	2.5	4.7
Median size	2	3
Minimum size	1	2
Maximum size	15	15

Source: HFCS Austria 2010, OeNB.

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

² The finite population correction accounts for the reduction in variance that occurs when sampling without replacement from a finite population. This type of sampling is used in the sample design for the HFCS in Austria.

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, but none of them have a value of 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, in addition, have larger maximum values than the final survey weights can be explained by the weight adjustments that were carried out: As some PSUs are removed from 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	1,586	1,429	169	9,054	3,773,956
1 st set of replicate weights	1,586	1,101	5	11,805	3,773,956
2 nd set of replicate weights	1,586	1,089	4	14,345	3,773,956
3 rd set of replicate weights	1,586	948	4	18,429	3,773,956
998 th set of replicate weights	1,586	1,174	3	22,191	3,773,956
999 th set of replicate weights	1,586	1,170	5	17,956	3,773,956
1,000 th set of replicate weights	1,586	1,122	3	14,139	3,773,956

Source: HFCS Austria 2010, 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 i.a. stratification, several stages of cluster-sampling and weighting adjustments, bias the variance estimators if data users ignore them.

While it is true that correctly calculating the standard errors by using replicate weights requires more computational power than analyzing the data without using replicate weights, in practice it is not necessary to use all 1,000 sets of replicate weights for variance estimation. Using fewer replicates, for example, yields faster (but less precise) variance estimations. The number of replicates used may depend on the type of estimator and on the size of the population surveyed. For instance, estimating the means of 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 have several special features that must be taken into account when analyzing these data. The data are multiply imputed, they have survey weights and replicate weights. The HFCS data also come in several files, accounting for the structure of the survey. These files are distinguished by the level of data (household vs. personal), by the number of implicates (i.e. each implicate is a separate file), and by type, i.e. whether the data were collected or constructed (derived variables, i.e. aggregate variables, and replicate weights vs. survey variables). This chapter¹ provides Stata^{®2} code that the user can employ step by step to account for all of these features.³ Part of the code was provided by Sébastien Pérez-Duarte⁴ of the ECB and is now available in a slightly altered and extended version. The ECB is expected to make available several program codes in spring 2013 as well. In this chapter, the program code for Stata[®] is given in the light blue-shaded area; it can be copied to the Stata[®] command window,⁵ and has to be run according to the steps that are laid out below (an alteration of the sequence may invalidate the code). 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 a way of merging the separate files, followed by the description of one possibility to set up the structure to use the imputations and survey information. Finally, some examples of simple estimation commands are provided in order to exemplify their use.

9.2 Merging the Data Files

The core HFCS data, which contain all internationally agreed variables, consist of the five multiple imputation samples or implicates at household level (files H1–H5), the corresponding sample at person level (files P1–P5) and the corresponding set of aggregate variables⁷ (files D1–D5). Before we start creating a new data set containing all these files, the user has to provide the computer-specific path to the data sets and the folder of the do-files that are used later on. The variables used for merging are the household identifier “sa0010”, the implicate number “im0100” and the country identifier “sa0100”.

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

¹ The authors do not make a judgement about which programs to use in a specific set-up. In particular if the size of the subsamples varies in each iteration, the estimation of discontinuous estimators does not comply with the assumptions underlying the findings on multiply imputed data in the literature (e.g. Little and Rubin, 2002). It is the responsibility of the user to check the validity and appropriateness of certain estimation commands in any given circumstances.

² The codes were written for Stata[®] version 12.1 and are not valid for previous Stata[®] versions. Comments are provided to facilitate the use of Stata[®] version 11.2.

³ Any changes to and improvements of the code are made available on an ongoing basis in the online appendix.

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 additionally provided do-files must be specified before execution. Due to the size and structure of the data and depending on the specification of the software and hardware, the execution of the do-file may take some time.

⁷ The ECB is expected to make available the definitions of the derived variables in spring 2013.

merging the resulting data set with the H-files. They are provided in *wide format*,⁸ i.e. one line of the data matrix includes information on one particular household and the information on each person within a household is put in a separate variable. Merged with the D-files, these M-files yield the entire HFCS data set in the file “hfcs.dta”.

```
*****
***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 "$hfcsdata\P%i.dta", clear
  drop id hid survey
  foreach var of varlist sa0010-fra0500 {
    local `var'lab: variable label `var'
  }
  reshape wide ra0?0* fra0?0* p* fp*, i(sa0010 sa0100) j( ra0010)
  foreach j of varlist ra* fra* p* fp* {
    local last2car=substr("`j'", `=length("`j")'-1', 1)
    local last1car=substr("`j'", length("`j"), 1)
    if "`last2car'=="1" {
      local firstcar=substr("`j'",1, `=length("`j")'-2')
      rename `j' `firstcar'_'last2car'`last1car'
      label variable `firstcar'_'last2car'`last1car' "`firstcar'lab' - `last2car'`last1car'"
    }
    else {
      local firstcar=substr("`j'",1, `=length("`j")'-1')
      rename `j' `firstcar'_'last1car'
      label variable `firstcar'_'last1car' "`firstcar'lab' - `last1car'"
    }
  }
  save "$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"
}

```

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.

```

*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 setting up the multiple imputed data and erased following this procedure.

9.3 Multiple Imputation

The next step is to import into Stata's[®] `mi` (i.e. `mi estimate` – commands for an appropriate use of the multiple imputation structure) both the original data and the imputed values. 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 macro “`IMPUTEDVARS`” in the program code below, which contains a string listing all imputed variables after execution of the corresponding loop. If registration was successful only a few variables (e.g. the implicate number “`im0100`”) and the flags should appear as “unregistered varying” when typing `mi varying` .

⁹ All missing values (including “Don’t know”, “No answer” and filter missings are set “.”. The flags make it possible to differentiate between these types of missing values (filter missing observations were flagged with a “0”). Flag variables have the same variable name, but start with an “f” preceding the variable name.

```

*****
***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"
  }
}

*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

*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

Once having declared the data as multiply imputed, we can proceed by designating the data as complex survey data. To this end we identify the variables that contain information about the survey design and specify the default method of 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 variables 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[®]. We write `– mi estimate: svy: –` followed by the estimation command in question. Stata[®] will then provide correctly calculated estimates and standard errors, taking into account both the multiple imputation framework and the replicate weights.¹⁰ The `– esampvaryok –` option can be useful in case that the sample size varies across imputations due to imputations.¹¹ Stata[®] versions below Stata[®] 12 do not allow the use of replicate weights together with multiply imputed data. Thus an underlying command¹² called `– u_mi_estimate –` has to be modified before the standard estimation commands can be used. In Stata[®] 12 the option `– vceok –` (used after the `– mi estimate –` command, e.g. “`mi estimate, vceok:...`”) can be used instead. It should be noted that in order to be able 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 (e.g. “`...svy, subpop(dummy)...`”).¹³ 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

*Modified Stata command, which should be run before the estimation commands for
versions of Stata previous to Stata 12 (always update to the most recent Stata
version). Furthermore, exclude in this case the vceok option in the estimations
below.
*do "$hfcsdofile\modified u_mi_estimate 11.2.do"

*Mean of current value of primary housing unit
mi estimate, esampvaryok vceok: svy: 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 (e.g. Little and Rubin, 2002) were derived assuming that the same set of observations is used in each imputed data set. They may thus not necessarily apply when the sets of observations used in the data analysis differ. This is why `– mi estimate –` errors out when this happens. In some cases, when the subsets used in each complete data analysis do not differ too much the conventional formulas may still be applicable. In this case, it is the choice of the user 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 the literature.

¹² The necessary alteration was written by the ECB and is provided as a separate do-file in the online appendix.

¹³ The use of an if-condition does not account for the uncertainty of the subsample size and therefore yields wrong variance estimators.

```

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

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

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

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

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

```

9.6 Using Additional Estimation Procedures

Calculating a median or another quantile requires a different Stata[®] package. This program is called – `medianize` – and provided by the ECB (the respective do-file can be found in the online appendix). This program must be used with caution as it has so far been tested only in limited environments. Furthermore, it uses the command – `tabstat` – and the analytical weights option of this command in Stata[®]. To our knowledge, however, there is presently no command to estimate discontinuous estimators (like median or other percentiles) other than the use of these ad-hoc measures.

```

*****
***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 that are necessary to use replicate weights together with multiple imputation in older Stata[®] versions and to estimate certain quantiles. The online appendix is intended to be updated on an ongoing basis to include program codes for various other HFCS-relevant topics. Every additional do-file will be supplemented with appropriate documentation.

References

- Albacete, N. 2012.** Multiple Imputation in the Austrian Household Survey on Housing Wealth. Oesterreichische Nationalbank. Working Paper 176.
- Banca d'Italia. 2012.** Sample Surveys – Household Income and Wealth in 2010. Supplements to the Statistical Bulletin. Year XXII – 25 January 2012, Number 6.
- Barceló, C. 2006.** Imputation of the 2002 Wave of the Spanish Survey of Household Finances (EFF). Documentos ocasionales 0603. Banco de España.
- Biemer, P. and S. Christ. 2008.** Chapter 17. Survey Weights. In: P. Levy and S. Lemeshow. Sampling of Populations: Methods and Applications. Fourth Edition. Hoboken, NJ: John Wiley & Sons.
- 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 to 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. Federal Reserve Bulletin 98(2). 1–80.
- Cameron, A. and P. Trivedi. 2005.** Microeconometrics: Methods and Applications. Cambridge University Press.
- ECB. 2011.** Core Output Variables Catalogue.
www.ecb.int/home/pdf/research/hfcn/core_output_variables.pdf?c6a87a29f0c1cdf4b92526aceef3efea (retrieved on January 22, 2013).
- Fessler, P., P. Mooslechner, M. Schürz and K. Wagner. 2009.** Housing Wealth of Austrian Households. Monetary Policy & the Economy Q2/09. OeNB.
- Heeringa, S. G., B. T. West and P. A. Berglund. 2010.** Applied Survey Data Analysis. Applied Survey Data Analysis. Chapman Hall / CRC Press: Boca Raton, FL.
- Iannacchione, V. G., J. G. Milne and R. E. Folsom. 1991.** Response Probability Weight Adjustments Using Logistic Regression. Proceedings of the American Statistical Associations, Section on Survey Methods. 637–642.
- 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 2005.
- Kennickell, A. B. 2011.** Look Again, Editing and Imputation of the SCF Panel Data. Prepared for the Joint Statistical Meeting. Miami, Florida. August 3, 2011.
- Kennickell, A. B. and D. McManus. 1993.** Sampling for Household Financial Characteristics Using Frame Information on Past Income. Proceedings of Survey Research Methods Section of the American Statistical Association. 88–97.
- Kish, L. 1995.** Survey Sampling. New York: Wiley.
- Kolenikov, S. 2010.** Resampling Variance Estimation for Complex Survey Data. The Stata Journal 10(2). 165–199.
- Kovar, J. G., J. N. K. Rao and C. F. J. Wu. 1988.** Bootstrap and Other Methods to Measure Errors in Survey Estimates. The Canadian Journal of Statistics 16.
- Levy, P. and S. Lemeshow. 2008.** Chapter 17, Variance Estimation in Complex Sample Surveys. In: Sampling of Populations: Methods and Applications. Fourth Edition. Hoboken, New York: John Wiley & Sons.
- Little, R. J. A. and D. B. Rubin. 2002.** Statistical Analysis with Missing Data. Wiley Series in Probability and Statistics. New York: Wiley Second Edition.

- Little, R. J. and S. Vartivarian. 2003.** On Weighting the Rates in Non-Response Weights. *Statistics in Medicine* 22(9). 1589–1599.
- Rao, J. N. K. and C. F. J. Wu. 1988.** Resampling Inference with Complex Survey Data. In: *Journal of the American Statistical Association* 83. 231–241.
- Rao, J. N. K., C. F. J. Wu and K. Yue. 1992.** Some Recent Work on Resampling Methods for Complex Surveys. *Survey Methodology* 18. 209–217.
- Royston, P. 2004.** Multiple Imputation of Missing Values. *Stata Journal* 4(3). 227–241.
- Schafer, J. L. and M. K. Olsen. 1998.** Multiple Imputation for Multivariate Missing-Data Problems: A Data Analyst's Perspective. *Multivariate Behavioral Research* 33. 545–571.
- Stata Library – Replicate Weights.** UCLA: Academic Technology Services, Statistical Consulting Group.
www.ats.ucla.edu/stat/stata/library/replicate_weights.htm (retrieved on January 22, 2013)
- Statistics Austria. 2011a.** Standard-Dokumentation Metainformationen (Definitionen, Erläuterungen, Methoden, Qualität) zu Mikrozensus ab 2004, Arbeitskräfte- und Wohnungserhebung, www.statistik.at/web_de/wcmsprod/groups/gd/documents/stddok/008863.pdf#pagemode=bookmarks (retrieved on January 22, 2013).
- Statistics Austria. 2011b.** Standard-Dokumentation Metainformationen (Definitionen, Erläuterungen, Methoden, Qualität) zu EU-SILC 2009.
www.statistik.at/web_de/wcmsprod/groups/gd/documents/stddok/059852.pdf#pagemode=bookmarks (retrieved on January 22, 2013).
- 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. *Statistics in Medicine* 18(6). 681–694.
- Van Buuren, S., J. P. Brand, C. G. Groothuis-Oudshoorn and D. B. Rubin. 2006.** Fully Conditional Specification in Multivariate Imputation. *Journal of Statistical Computation and Simulation* 76, 1049–1064(12).
- Vehovar, V. 1999.** Field substitution and unit nonresponse. *Journal of Official Statistics* 15. 335–350.

Monetary Policy & the Economy provides analyses and studies on central banking and economic policy topics and is published at quarterly intervals.

Publisher and editor

Oesterreichische Nationalbank
Otto-Wagner-Platz 3, 1090 Vienna, Austria
PO Box 61, 1011 Vienna, Austria
www.oenb.at
oenb.info@oenb.at
Phone (+43-1) 40420-6666
Fax (+43-1) 40420-6698

Editorial board

Peter Mooslechner, Ernest Gnan, Franz Nauschnigg, Doris Ritzberger-Grünwald,
Martin Summer

Managing editor

Karin Wagner

Editing

Rita Schwarz

Translations

Dagmar Dichtl, Jennifer Gredler, Rena Mühlendorf, Ingeborg Schuch, Susanne Steinacher

Design

Communications Division

Layout and typesetting

Walter Grosser, Birgit Vogt

Printing and production

Web and Printing Services

DVR 0031577

© Oesterreichische Nationalbank, 2013. All rights reserved.

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

Printed according to the Austrian Ecolabel guideline for printed matter.

