



universität
wien

DISSERTATION

Titel der Dissertation

“Essays in Economic Survey Analysis”

Verfasser

Mag. Pirmin Fessler

angestrebter akademischer Grad

Doctor of Philosophy (PhD)

Wien, 2012

Studienkennzahl lt. Studienblatt:
Dissertationsgebiet lt. Studienblatt:
Betreuer:

A 094 140
Volkswirtschaftslehre (Economics)
Univ. -Prof. Dipl. -Ing. Dr. Robert M. Kunst

Contents

0.1	Introduction	III
1	The Gendered Aspects of the Intergenerational Persistence of Educational Attainment in Austria	1
1.1	Introduction	1
1.2	Literature Overview	6
1.3	Empirical Analysis	13
1.3.1	Data	13
1.3.2	Univariate Analysis	18
1.3.3	Markovian Approach	23
1.3.4	Multivariate Analysis - Ordered Logit	29
1.4	Conclusions	34
1.5	Literature	37
2	Assets and Liabilities of Austrian Households	43
2.1	Introduction	43
2.2	The macro level	45
2.3	Debtors in Austria	47
2.3.1	Real estate debt	49
2.3.2	Consumer debt	52
2.4	International Comparison of selected OECD countries	54
2.4.1	Empirical Results for Selected OECD Countries	55
2.5	Empirical Results for Austria	58
2.6	Conclusions	65
2.7	Literature	68
3	The Role of Data Production in Survey Analysis	71
3.1	Introduction	71
3.2	How survey production can drive the results of survey analysis	72
3.2.1	Sampling frame and sampling design	73
3.2.2	Survey mode	78
3.2.3	Interviewing	85
3.2.4	Paradata	86
3.2.5	Weights	87

3.2.6	Imputation	94
3.3	Final Remarks	98
3.4	Literature	100
4	A Note on the Effects of Higher Education on Earnings	103
4.1	Introduction	103
4.2	Empirical Strategy	105
4.2.1	Potential Outcome Approach	105
4.2.2	Average Treatment Effects on the Treated	107
4.3	The data	109
4.4	Results	112
4.5	Conclusion	117
4.6	Literature	120
A	Abstracts	123
A.1	English Abstract	123
A.1.1	The Gendered Aspects of the Intergenerational Persistence of Educational Attainment in Austria	123
A.1.2	Assets and Liabilities of Austrian Households	124
A.1.3	The Role of Data Production in Survey Analysis	124
A.1.4	A Note on the Effects of Higher Education on Earnings	124
A.2	German Abstract	126
A.2.1	The Gendered Aspects of the Intergenerational Persistence of Educational Attainment in Austria	126
A.2.2	Assets and Liabilities of Austrian Households	126
A.2.3	The Role of Data Production in Survey Analysis	127
A.2.4	A Note on the Effects of Higher Education on Earnings	127
B	Curriculum Vitae	129
C	Acknowledgements	137

0.1 Introduction

This dissertation consists of four chapters which are all based on or deal with “Economic Survey Analysis”. The economic analysis of survey data has a long tradition in the field of economic analysis. There has been a strong improvement in the process of data gathering and as a result, more and more high quality datasets are available for scientific analysis. In Chapter 3 on the “Role of Data Production in Survey Analysis” we discuss important issues regarding the interplay of the process of data production and the process of data analysis in the case of survey data. Chapters 1, 2 and 4 are all based on household survey data and bring forth arguments on principal points in the economic literature which are mostly overlooked and deserve closer attention. Furthermore, all three of these chapters provide the first analysis using Austrian data to analyse the respective topic.

Chapter 1 examines the differences of same-gender versus cross-gender patterns in the process of the intergenerational transmission of educational attainment in Austria. This is the first analysis documenting the important role of gender patterns in the intergenerational transmission of educational attainment. In Chapter 2 we investigate the joint distribution of assets and liabilities of Austrian households. This is the first analysis dealing explicitly with the relationship between assets and liabilities on the basis of a micro-level dataset for Austria. While economic theory would suggest that people with high amounts of debt tend to hold zero or few liquid assets empirical analysis show the strong positive association between assets and liabilities. Finally in chapter three we estimate average treatment effects on the treated to analyse the effects of higher education on earnings and show how returns on education in Austria are usually overestimated when ignoring selection bias based on parental education. More detailed abstracts can be found in the appendix.

Chapter 1

The Gendered Aspects of the Intergenerational Persistence of Educational Attainment in Austria

This Chapter is based on joint work with Alyssa Schneebaum and published in a slightly different version in *Feminist Economics*, Fessler, P. and Schneebaum, A. (2012): Gender and Educational Attainment Across Generations in Austria, *Feminist Economics*, 18:1, 161-188.¹

1.1 Introduction

Gender-based differentials in educational attainment are an important part of the history of women's oppression and an important aspect of the analysis of women's rights and opportunities. In this paper, we examine the stylized fact that descendants of parents with higher levels of education are more likely to achieve higher levels of education

¹A related paper - using the same dataset - which focuses more on delivering internationally comparable results for Austria instead of the gender issue is also already published (Fessler, P., Mooslechner, P., and Schürz, M. 2012)

than descendants whose parents have less education, and we analyze the results by the genders of the parents and children involved. Our gendered analysis is meant to show that gender is an important variable in considering all aspects of one's education, playing a part in even the basic fact that there is intergenerational persistence in educational attainment.

In many instances, advantages and disadvantages are passed from one generation to the next. A society that is characterized by a high degree of transmission of social status may have problems in claiming meritocratic ideals; the same constraint is true regarding gender-based equality if the educational transmission is characterized by differences for men and women. Our main focus lies on examining the differences in the educational transmission process between parents and children of the same gender compared to parents and children of different genders.

We therefore test the following questions:

Q1. Is there an intergenerational persistence in educational outcomes, i.e. is the education of parents and descendants positively correlated?

Q2. What is the relevance of gender in determining intergenerational persistence? Are there any differences between same gender parent-children pairs and cross-gender parent-children pairs?

Q3. Has the strength of the persistence varied over time, and if so, why?

We use a Markovian approach, along with univariate and multivariate econometric techniques, to test these questions. The range of methods employed allows us to check the robustness of our results. Due to the absence of long panel data series for

Austria, we use the Household Survey on Housing Wealth (HSHW), a cross sectional survey which collects information on the respondents' education along with data on the educational level of their parents. We discuss the data in greater detail below.

The main question in this paper focuses on the gendered aspects of the intergenerational transmission of education. As such, we pay careful attention to the trends in educational attainment by gender. Before the 1970s, women had much lower rates of education than their male peers. However, in the 1970s, the Austrian government introduced several policies to increase educational attainment in general, but particularly for children from low-income backgrounds. Given that females had significantly lower educational attainment on average, they benefited more than males from this policy change, as it increased educational access to previously constrained households and families who presumably sacrificed their daughter's educational opportunities more often than they did their son's. Especially under the assumption that marginal returns to education decrease with years of education, we would expect to see women experiencing a stronger benefit from an expanse in public education funding, since females had lower levels of education. Indeed, we can see that this policy helped females boost their educational accomplishments. Moreover, there were other policy measures introduced in the 1970s that promoted women's rights in general, alongside the formation of the feminist movement in Austria, which may have also had an impact on gender roles, educational accomplishments, and educational transmission. We address the issue of the educational reform and changes in women's rights, and how those institutional changes could influence the findings of our study, in more detail below.

In this study, we concentrate in the intergenerational correlation of education (Hertz et al., 2007; Mulligan, 1999). While there is a strand of literature aimed at identifying total causal effects of parents' education on the educational attainment of their children via twin datasets (Behrman and Rosenzweig, 2002), adoptee datasets (Plug, 2004), or

school reforms (Black et al., 2005) to control for parents' unobserved endowments, we instead look at correlation, not causation. Some of the recent literature on the intergenerational transmission of various characteristics - such as education - focuses on causal effects and we find that discussion important, so we highlight it below. However, the causal approach is not the one we take in our empirical analysis. Generally speaking, there are numerous factors that shape intergenerational mobility, described as follows. Household income is a critical component of intergenerational transmission. Income poverty is related to negative outcomes such as poor health and low levels of nutrition and housing. Such conditions negatively affect the future life-chances of children. Because it cumulates with other forms of disadvantages (such as a lack of access to educational resources), low income mobility is especially harmful to the development of children's opportunities. The income of the household into which a child is born will greatly affect the child's access to resources and therefore class mobility, making income class perpetual over generations. Children coming from high-income households have more material resources to aid their social and intellectual development. Children in these households may also positively benefit from a transmission of desirable verbal skills and non-cognitive abilities such as motivation, socioemotional regulation, time preference, personality factors and the ability to work with others (as summarized by Heckman, 2008) while children in low-income households may receive lower levels of those skills via intergenerational transmission. Children from high-income homes also benefit from access to social networks that help their professional development. On the other hand, there are particular skills that children in a working class family might learn while their more economically well-off peers are excluded from that information. However, the extra skills (say, for example, handicraft work) are likely to be useful mainly, although of course not exclusively, for low-income jobs since it was workers in those trades that taught it to the children, which furthers the intergenerational depen-

dence of income class transmission.

Household wealth is also an important component of intergenerational mobility. Wealth constrained parents cannot invest as much in resources (including education) for their children as can wealthy parents. Borrowing against future earnings is difficult, and liquidity constraints affect investment in human capital (Becker and Tomes, 1979). Wealthy parents can pass gifts and bequests on to their children, deepening the relationship between one's own situation and his or her family background. Such wealth transfers increase the asset holdings of only a portion of children and deteriorate the principle of equal opportunity.

Educational attainment is significantly correlated across generations. Educational traits persist between generations in all OECD countries, and the OECD claims that parental education is by far the most important background characteristic in determining a child's level of education (see OECD 2008, p. 216). Belzil and Hansen (2003) argue that household background variables - particularly parents' education - account for 68% of the explained cross sectional variations in schooling. The persistence of educational attainment between parents and children plays a major role in limiting intergenerational mobility in general, especially because one's level of education so heavily influences one's life position.

Genetic factors might also matter when thinking about why certain outcomes persist across generations. Children may inherit genetically based behavioral characteristics from their parents. Herrnstein and Murray's 1994 book *The Bell Curve*, for example, argues that differences in outcomes are based primarily on genetic factors. However, the contribution of genetic factors in studies of intergenerational mobility remains rather unclear, and the study presented in *The Bell Curve* has been strongly disputed (Hansen et al. (2004) provide one case for why that study was not relevant). Bowles et al. (2005) find that little intergenerational inequality is due to parents passing IQ on to their chil-

dren. Heckman (2008) shows that in the US, increased monetary investments during early childhood and adolescence greatly increase high school graduation rates and college enrollment (p. 70), suggesting that social influences are far more important than biological ones in determining a child's educational success.

Many additional factors are interwoven in the process of intergenerational transmission of inequality. Aside from wealth, income, parents' educational attainment, and genetics, a child's social environment and his or her household structure is closely related to this phenomenon. A thorough empirical assessment of every important characteristic would require a survey containing data on a wide range of individual and social characteristics of parents and children. Like most surveys, the HSHW does not include such extensive information. However, we do have data on parent's education, and as numerous studies have shown, this is a strong indicator for intergenerational inequality because it is closely related to income and wealth. We therefore focus our attention to the strength of educational persistence as an important measure of intergenerational (im)mobility, by concentrating on the correlation between two generations.

1.2 Literature Overview

The theoretical background of most empirical models on intergenerational transmission is the Becker-Tomes (1979, 1986) model, which is itself related to Francis Galton's 1877 work (Mulligan, 1999). Both are heavily discussed in the literature on intergenerational transmission and transfers (Bowles and Gintis, 2001). Intergenerational correlation of income, wealth, consumption, and education is well documented in a tremendous number of empirical studies, as shown by Mulligan (1999, Table 1). Of course, the reasons for the intergenerational correlation of various characteristics, including educational

attainment, may be multidimensional. The literature identifies genetic transmission, parental care, parental abilities, parental role modeling, family income and wealth, pre-school and school facilities, and one's out of school environment as the most prominent factors in determining a child's educational achievements. A central discussion is the 'nature versus nurture' question, i.e. whether the high correlation between parents' and descendants' characteristics is mainly due to the genetic transmission of ability or to the social environment of the descendants. No consensus exists on this question, but most researchers agree that the answer lies mainly in the child's given social environment, especially the resources endowed to the parents (Checchi et al., 2008; Heckman, 2008). Speaking specifically of educational attainment, the literature at hand shows that parents' education is the most important factor explaining the educational attainment of children (Haveman and Wolfe, 1995). The main difficulty in testing the ubiquity of this result is the fact that many datasets do not provide adequate information; when we do have a dataset that contains educational information of both the parents and descendants, it often lacks many human capital variables which we would use to control for abilities of descendants to estimate direct causal effects in educational transmission (Dardanoni et al., 2008). Nevertheless, it remains unclear if abilities, even tested at a young age, are not already formed by the social environment and especially parental education. With respect to strategies to control for parents' unobserved endowments in order to identify total causal effects, it remains unclear if parents' child rearing abilities of twins (parent-generation) are identical (Behrman and Rosenzweig, 2002) or, in the case of adoptee datasets, if the process of adoption is random or if there is some selection which would be comparable to inheritable abilities (see Plug, 2004). Furthermore, different approaches to control for parental endowment can lead to different and contradictory results (Holmlund et al. 2008). Given the difficulty in measuring causation within the intergenerational persistence framework, analyzing intergenerational

correlation instead seems to be valuable, even in the face of the possibility that the relationship between the education level of parents and their children could be overestimated. Therefore, our estimates are to be interpreted as correlations, not as direct or total causal effects.

There is also a substantial body of literature in sociology dealing with inheritance of inequalities across generations for different countries (D'Addio, 2007). Sociologists have been mainly concerned with the intergenerational transmission of attitudes and values. Loehlin (2005) estimates the correlations between parents and their children's personality traits, attitudes, values and interests across various family types and finds a significant positive correlation for attitudes, values and interests. Osborne-Groves (2005) reviews estimates of intergenerational correlation of personality traits in 'ordinary families' based on several studies and argues that personality traits are both persistent across generations and relatively stable over time. Bourdieu (1984) emphasizes the relevance of economic, cultural and social capital for the reproduction of class inequality, but in his framework, individuals are not assumed to follow self-interested strategies in a rational agent sense as they are in economics but instead their habitus², a basic concept of Bourdieu's sociological theory.

With respect to the particular institutional arrangements of Austrian society and education, there are several studies that document the continuing educational expansion in Austria. Since the Monarchy, there has been a continuing positive trend - with some phases of stagnation and an interruption during the Second World War - towards higher education. Beginning with the cohorts of the mid to late 1950s who started primary school in the early to mid 1960s, we see a constant educational expansion starting with

²In Bourdieu's work, habitus is a system of dispositions (perception, thought and action). The individual agent develops these dispositions in response to the determining structures (such as class, family, and education) and external conditions (fields) they encounter. They are therefore neither wholly voluntary nor wholly involuntary.

around 6% for Matura³ and higher education and reaching 25% and more in the 1990s (Steiner, 1998). There are studies dealing with social selection in the Austrian schooling system based on school statistics. In general, Austrian school statistics do not allow for analyses of the influence of social background (e.g. education of parents), except for university students (Landler, 1997). Other studies, as Bacher (2003) shows, use survey data of adolescents living with their parents to analyze social selection in secondary schools. Both approaches allow for a detailed analysis of their subjects but are limited in two ways: they do not include the whole population and they do not include the entire educational system.

The study that is most closely related to our analysis is that of Spielauer (2004), which analyses the intergenerational educational transmission within families and finds that there is considerable intergenerational persistence of educational attainment, a result which counteracts the strong educational expansion over the past several decades. One drawback of Spielauer's approach is that he uses the maximum of both fathers' and mothers' education to analyze the effect of parents' education on the child's probability of attending different school types and having various educational outcomes. Therefore, it is not possible to compare either the influence of mothers' education with the influence of fathers' education, or same-gender parent-descendant relationships with cross-gender parent-descendant relationships. The main aim of Spielauer's study is to analyze school choices and educational paths, and to compute a projection of the distribution of educational attainment of the society. In contrast, our focus lies on the educational outcome and how it is transmitted from one generation to the next. Our study includes, as in Spielauer's sample, a reasonably good representation of the whole population and therefore the whole educational system, and includes an important analysis by gender of parent and child. Furthermore, the selected empirical strategy allows

³Matura is more or less comparable to high-school (upper secondary education)

- at least partly - for international comparisons.

We know of only limited literature that looks specifically at the gendered aspects (with respect to parents and descendants) of the transmission of educational attainment, and none that emphasize the importance of gender in the topic of intergenerational persistence. Dardanoni et al. (2008) found a significant father to son causal effect of parental schooling and an insignificant weak effect of mothers on daughters. Dearden et al. (1997) looked at earnings and educational transmission of both parents to their children in the UK, and found that fathers' education level has a stronger effect on the outcomes of their sons, while that of mothers has a stronger effect on the outcomes of their daughters. Moen et al. (1997) investigated the transmission of gender roles from mothers to daughters, and find a positive correlation that has decreased since the 1960s. This finding - as well as those in the sociological literature (see above: Bourdieu, 1984; Loehlin, 2005; Osborne Groves, 2005) - could be relevant in our case; if a mother has a certain view about women's education in general, she may pass that ideology along to her daughter. We hypothesize that the relationship between parents and children of the same gender will be stronger than the cross-gender relationship, in part because of this preservation of gender role ideology between parents and their children. Not only do parents teach their ideas about gender roles (what is expected of men and women) to their children, but their own educational experiences could have affected their notions about gender expectations, strengthening the connection between the educational attainment of parents and their same-gender children.

We further expect - in line with Spielauer (2004) - that both males and females will have had more mobility over time which maybe partly due to legislative changes and investments in the Austrian education system, and we expect females to have experienced greater educational gains over time that could be partly due to the feminist movement, which has changed gender ideology and encouraged females to pursue education. In the

few years from 1970 onwards, the Austrian government changed the educational system dramatically.

In June 1971 the parliament passed the 4th school organization amendment. The center of the amendment was the abolition of the entrance exam into grammar school (Allgemeine Höhere Schule). However, the division between secondary modern school (Hauptschule) and grammar school (Allgemeine Höhere Schule) at the age of 10 is still in place and might be relevant for a relatively low level of educational mobility (see Spielauer 2004) up to date. In 1971, free public transportation for pupils was introduced in order to open access to educational channels that lead with greater probability to higher education for lower social classes (grammar schools are concentrated in cities or larger villages whereas modern secondary schools are wider spread). In the autumn of 1972, pupils received free school books for the first time. Two years later, on September 1, 1974, a new school organizational law (Schulorganisationsgesetz) brought some limitation of the rights of the teachers and an improvement of the rights of the pupils. Concerning university education, a law passed in 1972 abolished all costs for studying (at least for Austrian citizens and citizens from developing countries) but also heavily increased scholarship funding. Along with those changes, the structure of university organization was notably democratized. The new University Organization Law (Universitätsorganisationsgesetz) brought students and university assistants a say in the university organization. This new law was brought into effect in April 1975, against the vote of the more conservative parties (Österreichische Volkspartei and Freiheitliche Partei Österreichs). It seems that women especially benefited from these policy measures, perhaps because the new policies were particularly helpful for low-income households and those that spent less on education in general. Given the lower average educational attainment of women during and before the 1970s, it is clear that the parents of this generation invested more in the education of their sons than in

their daughters. We therefore speculate that the educational reforms allowed women to receive higher education without a monetary strain on their families. However, although it appears clear that the reforms had an effect on women's educational attainment, the magnitude of the causal effect of the reform cannot be measured by our approach (nor with our limited dataset) and we suggest further research on the effect of these policies. Alongside political reforms in the educational system, the feminist movement may have played an important role in promoting women's education and changes in persistence of educational traits between generations, by challenging traditional gender roles. In this respect the first half of the 1970s marked a turning point: the first autonomous feminist movement AUF (Aktion Unabhängiger Frauen) formed in the early 1970s in Vienna and founded their own magazine in 1974. The formation took place in the struggle for the liberalization of anti-abortion laws, where the autonomous feminist movement and women from the communist party and social democratic party worked together in order to legalize abortion. In 1975 the so called Fristenlösung, a law which legalized abortion up to a certain age of the fetus, was installed. In 1973 the tax system was changed towards a individual concept, whereas before it was a family based concept, where the woman's income (if married) was always part of the total family income and therefore in terms of taxes never recognized as an individual. In 1975 the Familienrechtsreform (Reform of Family Rights) passed the Parliament, which included for the first time full legal equal rights for men and women. Furthermore, in 1975 a law was passed which allowed for real estate ownership to be split among more than just one person, which led to many more women legally owning homes. Beginning in 1976 (with the passage of a new law called Unterhaltsvorschußgesetz), the state makes alimony payments when a father eludes to pay, and the state enforces the debt of the father, which makes divorce much easier for women. Furthermore, the marriage law was significantly changed in 1976, when equal legal rights for both partners were introduced. Since 1977, women

are allowed to keep their name after marriage or combine the names of the partners. Since 1978 both parents have the same child custody rights. In 1979, a law for the enforcement of equal pay for equal work was installed - discrimination in collective bargaining agreements are forbidden. Also in 1979, women's policy got institutionalized by creating the first state secretary for women's rights concerns (Staatssekretariat für allgemeine Frauenfragen). In the following years, several more legal measures in order to foster equal rights were introduced. Further, in the cultural and public spheres, the feminist movement played an increasingly important role from the 1970s onwards, resulting in a diversity of organizations, projects, initiatives, political activist groups, women's centers, magazines, and libraries dealing with feminist issues. In the academic sphere as well, change is coming, although complete equality is far from having been achieved. The higher the positions in the academic sphere, the slower the change occurs: while in 1977 the share of female university assistants was around 22%, it increased to around 35% in 2007. Concerning professorships, the share increased from around 2% in 1977 to around 15% in 2007 (BMWF 2007). We assert that these vast changes to women's rights and autonomy had some positive effect on women's opportunity and motivation to obtain an education and helped to close the gender gap in educational attainment in Austria.

1.3 Empirical Analysis

1.3.1 Data

To analyze intergenerational processes, one needs data incorporating information on at least two generations, usually a descendant and his or her parents. There are few Austrian datasets containing this information for a representative sample of descendants. The present study is based on empirical data collected for the Household Survey

on Housing Wealth 2008 (HSHW 2008) of the Austrian Central Bank (Oesterreichische Nationalbank), which was conducted as a pilot project for a planned comprehensive Eurosystem household survey⁴. It is a representative household survey investigating the housing wealth of Austrian households, with a 65.1% response rate. The respondents were either the owners or tenants of the respective household's real estate at the time of the interview. The survey focuses on the ownership of the respective house/apartment and of additional real estate belonging to any of the household members as well as on the related liabilities owed by the household. Furthermore, detailed socio-economic characteristics and data concerning intergenerational transfers in connection with housing wealth are compiled from these data. The questionnaire contains a total of 168 questions, 28 of which are related to socio-economic characteristics (additionally, 8 questions had to be answered personally by the interviewers themselves). The HSHW 2008 uses a stratified multistage cluster address random sample⁵. It was carried out using a computer-assisted personal interviewing (CAPI) method, which allows for immediate plausibility checks during the course of the interview, thus making it possible to correct for inconsistencies right away. The survey was conducted in January, February and March of 2008 with fieldwork taking approximately nine weeks. Comprehensive follow-up research was carried out until September 2008. The interviewee was asked to state the educational level of his or her mother and father. Out of the 2081 total observations, we use only 1892 observations for respondents aged 25 and older who have no missing information regarding their own education or the education of their mothers or fathers⁶. To check the validity of our data we first compare our variable of interest

⁴The HSHW 2008 was conducted by the Institute for Empirical Social Studies (IFES). For Information on the planned Eurosystem Household Survey on Finance and Consumption (HFCS) see http://www.ecb.int/home/html/researcher_hfcn.en.html.

⁵See Wagner and Zottel (2009) for further details.

⁶For 50 observations fathers education is missing and for 37 observations mothers education is missing. Age is missing for one descendant and 127 descendants are younger than 25. Due to overlaps of some of the observations we have to exclude only 189 observations.

- educational attainment of the population - with another reliable source. Table 1.1 compares the descendants' educational distribution of males and females with actual population statistics of Statistics Austria⁷. The distributions are reasonably close to each other; for each of the five categories reported the ordering of males and females is the same. The differences, however small, can be explained by several factors. First, the data of Statistics Austria is based on a survey as well, and therefore both surveys deliver just an estimate of the true distribution. Second, we use only the observations of descendants for whom we have an age and who are not missing information about their parents' education. Third, in our selected sample all people above the age of 25 are included, whereas Statistics Austria just reports Education for the Population between 25 and 64. If we restrict our sample to the people of age 25 to 64, our distributions are even closer to those of Statistics Austria, but for the purpose of our analysis, restricting our sample to people between those ages is not helpful and it only results in a loss of observations.

Table 1.1: Comparison of Educational Distribution of Descendants in the Survey

		HSHW 2008		Statistics Austria 2008	
		Male %	Female %	Male %	Female %
Low Education	max. Compulsory school level	13.7	19.7	12.5	22.3
	Apprenticeship or vocational school degree	45.5	37.8	51.4	30.3
Medium Education	Medium-level or technical school	13.5	16.8	8.9	18.7
	Matura and higher level vocational school	13.9	16.0	15.7	18.9
High Education	University, Fachhochschule	13.4	9.7	11.5	9.8

To further evaluate the validity of our data concerning educational attainment in

⁷Statistics Austria, Microcensus 2008.

combination with the age structure of descendants and their parents, we calculate descriptive statistics for different age groups of descendants. For reasons of comparability, the goal of our analysis, and the limited number of observations, we form four age groups and report educational attainment (for now) as the share of people with Matura and higher education - as in Steiner (1998, Table 1.2), for example. We formed age groups as follows: those born before 1945, which was the end of the Second World War; those born between 1945 and 1959, which includes the cohorts before or at the edge of the start of the educational expansion (see e.g. Steiner 1998); those born between 1960-1969, the first category to clearly benefit from school reforms in the 1960s and university reforms in the 1970s; and those born between 1970-1983, the second category to benefit from the educational expansion. The different time spans for our age categories do not matter for the following analysis and are just established for interpretational reasons (i.e. two before and two after the start of the strong educational expansion), reasons of inclusion of the whole sample (e.g. 1970-1983 instead of 1970-1979 to include all descendants aged 25 or older), and reasons of sample size (which is, as an example, too small to produce meaningful age categories of 5-year cohorts).

The average age of descendants in our selected sample is 50 years old and it is the same for females and males. In the oldest age category (those born before 1945), the average female is slightly older than her male counterpart, which might be due to the higher life expectancy for females. Further, the age structure of the parents seems to be reasonable: for those in the oldest age category, the average age difference between descendants and their fathers is about 26-27 years while it is about 23-25 years between descendants and their mothers. The difference in age between generations increases for the younger age categories. The average difference between the age of the fathers and the mothers is about 2-4 years for all age categories. Concerning educational attainment, those in the descendant population generally have higher levels of education

Table 1.2: Descriptive Statistics on Age and Education of Descendants and their Parents

n=900		Male descendants			Fathers			Mothers		
Age-Category	Mean Age	Share Matura and higher	Mean (Hypothetical) Age	Share Matura and higher	Mean (Hypothetical) Age	Share Matura and higher	Mean (Hypothetical) Age	Share Matura and higher	Mean (Hypothetical) Age	Share Matura and higher
born before 1945	72.3	0.21	98.9	0.09	97.0	0.06				
born 1945-1959	55.1	0.26	86.3	0.14	83.0	0.06				
born 1960-1969	43.5	0.30	73.7	0.15	70.1	0.06				
born 1970-1983	32.1	0.32	61.7	0.19	58.1	0.13				
Total	49.5	0.27	77.5	0.15	74.8	0.08				

n=992		Female descendants			Fathers			Mothers		
Age-Category	Mean Age	Share Matura and higher	Mean (Hypothetical) Age	Share Matura and higher	Mean (Hypothetical) Age	Share Matura and higher	Mean (Hypothetical) Age	Share Matura and higher	Mean (Hypothetical) Age	Share Matura and higher
born before 1945	73.8	0.12	99.5	0.08	97.2	0.03				
born 1945-1959	55.1	0.26	86.7	0.13	83.5	0.06				
born 1960-1969	44.0	0.31	74.6	0.21	70.9	0.08				
born 1970-1983	32.0	0.31	61.1	0.15	58.2	0.10				
Total	50.0	0.26	77.5	0.15	75.0	0.07				

than their parents. Furthermore, fathers are generally more educated than mothers and the steady expansion of educational attainment is clearly visible for both genders - but stronger for females. In all these aspects, the data tie in with previous studies about the educational expansion in Austria (e.g. Steiner 1998).

1.3.2 Univariate Analysis

As discussed above, the education level of a child depends on many factors. However, most studies exploring the intergenerational transmission of education concentrate on the correlation between parents' and descendants' educational attainment. This is because most data do not include good measures of social environment, parental care, or wealth. As a result, most studies have to assume that educational achievement includes the other aspects, at least in part. The general functional form of the following estimations will therefore be $E_i^d = E_i^d(E_i^f, E_i^m, C_i^d)$ for $i = 1, 2, \dots, N$, where E_i^d, E_i^f, E_i^m describes the individual educational attainment of an individual from the descendant pool and her father's or mother's education respectively, and C_i^d are additional characteristics of an individual belonging to the descendant population.

Univariate Analysis - OLS and Correlation. In order to be able to compare our results with those of other countries, we use univariate methods. Such techniques have been heavily used to analyze intergenerational transmission of educational attainment for a large number of countries (Chevalier et al., 2003; Hertz et al. 2007). Following the approach by Checchi et al. (2008), we estimate OLS regressions of the form

$$E_i^d = \alpha + \beta E_i^p + \varepsilon_i \quad \text{for } i = 1, 2, \dots, N, \quad (1.3.1)$$

where $p = f$ (father) in the first estimation and $p = m$ (mother) in the second estimation. Furthermore ε_i is assumed to be a normally distributed error term with

zero mean and σ^2 variance. The OLS estimate for each regression is

$$\hat{\beta} = \frac{\sigma_{dp}}{\sigma_p^2} = \rho_{dp} \frac{\sigma_d}{\sigma_p},$$

where $\sigma_d, \sigma_f, \sigma_m$ are the standard deviations of education of the relevant populations and ρ_{dp} is the correlation coefficient between descendants and fathers ($p = f$) or mothers ($p = m$) education. A decreasing $\hat{\beta}$ over time represents greater descendant independence in their educational outcomes vis-à-vis their parents. To ensure that a possible decrease or increase in $\hat{\beta}$ is not due solely to an evolution of the distributions of the educational attainments - namely the term $\frac{\sigma_d}{\sigma_p}$ - one can normalize the individual educational attainment variables by the corresponding standard deviations. This process provides an intuitive correlative interpretation, and leads to

$$\frac{E_i^d}{\sigma_d} = \alpha + \gamma \frac{E_i^p}{\sigma_p} + \varepsilon_i \text{ for } i = 1, 2, \dots, N \quad (1.3.2)$$

where the changes in γ over the separately estimated age categories can be interpreted as the evolution of the correlation between parents' and descendants' education levels. Note that for this exercise we had to transfer the categorical variables into statutory schooling years, i.e. the years which are at least necessary to complete a certain educational degree⁸.

Furthermore, as our data do not allow for instrumental variable estimation, the interpretation of the level of the estimates may be biased in relation to a causal interpretation due to the lack of controls for parental care, parental ability, social environment and other related but missing factors. Regardless, a causal interpretation of the changes

⁸In doing so we use all the categorical information available and replace them with appropriate statutory schooling years: maximum compulsory school=9, apprenticeship and vocational school=10, medium technical school=11, Matura and higher vocational school=12.5, University and Fachhochschule=16. Due to the complex educational system it is not unambiguously clear which values would be the most appropriate. However we tested a set of reasonable values and the results are fairly robust.

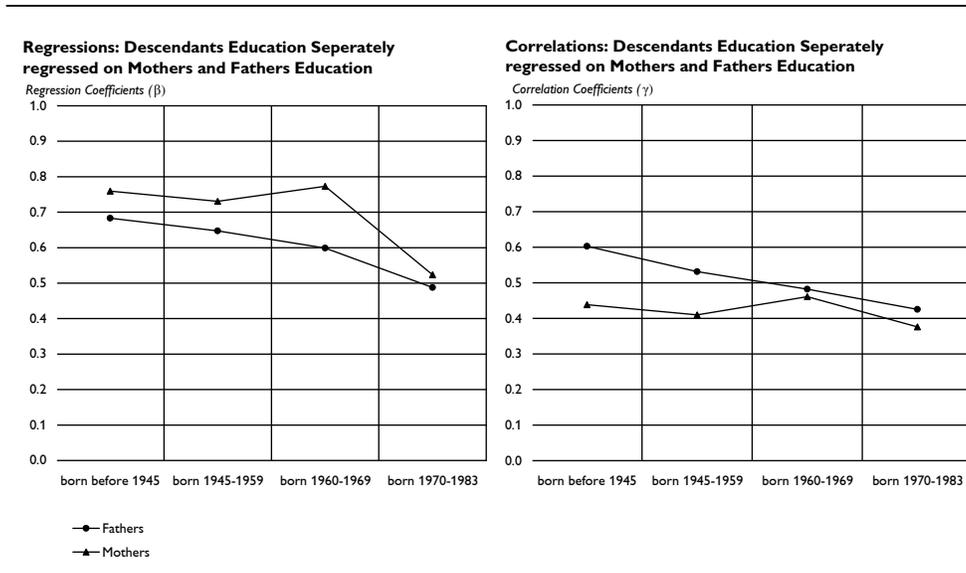


Figure 1.1: Educational Regression- and Correlation Coefficients over Age Categories

over time would be valid under the assumption that the influences of the possible biasing factors are time invariant. Of course this assumption is highly speculative and therefore a correlative interpretation - as we have pointed out above - is in order. Figure 1.1 shows the estimated coefficients of intergenerational educational transmission of Models 1 and 2 with (i) fathers' education level as the independent variable and (ii) mothers' education level as the independent variable .

The dependence of the educational attainment of the descendants on their parents' educational accomplishments decreased significantly over time. In other words, children starting school more recently have had greater educational mobility. The fact that we find higher β coefficients for the mothers' regressions than for the fathers' regressions but the reverse for the γ coefficients shows that a large portion of the β coefficient is due

to differences in the distributions⁹. Overall, fathers' as well as mothers' education correlates significantly - in terms of β and γ coefficients, which all are significantly different from zero at a 1% significance level - with their descendants' education. The correlation has decreased over time, with the sole exception of the coefficients for mothers, which increased slightly from the age category 1945-1959 to the age category 1960-1969 but decreased sharply to the age category 1970-1983. Disregarding distributional differences of the population and their changes over time the correlation between fathers and descendants is higher than that between mothers and descendants. Furthermore, the β coefficients decreased more dramatically than the γ coefficients, implying that we would overstate the decrease in the level of persistence if we did not control for the distributional differences over age categories resulting from the educational expansion in that period. This result offers further justification as to why it is appropriate to split the sample in 4 age categories and estimate the coefficients on the different age subsets, instead of using a single estimation with dummies for age categories. Furthermore, for an analysis regarding the evolution of intergenerational transmission over time, it is important to avoid the different sizes of birth-cohorts, which would have different weights for the overall results. Of course in that sense it would be better to have even more subsets of age categories, but only if the sample size were large enough.

With respect to the trend and magnitude of the evolution of the coefficients, our results are in line with those of Checchi et al. (2008) for Italy. Compared to measures estimated¹⁰ by Hertz et al. (2007) our results seem to be quite reasonable as well. Their estimated β coefficient for Italy is 0.67, 0.58 for Sweden and the Netherlands, 0.54 for

⁹The starting level of the standard deviations of statutory schooling years in the oldest age cohort is clearly lower for the mothers than for the fathers, and the standard deviation of all (mothers, fathers, and descendants) is rising over time. The standard deviation of the mothers' population rises strongly but remains the lowest in all four subsets. The different evolution of standard deviations explains the evolution of the differences between the β and γ coefficients.

¹⁰For their estimation, Hertz et al. (2007) used the average years of school of fathers and mothers and calculated overall coefficients by averaging cohort coefficients.

Slovenia, 0.48 for Finland, and 0.46 for the USA. The correlation estimate γ (disregarding distributional changes) is 0.54 for Italy, 0.52 for Slovenia, 0.46 for the USA, 0.40 for Sweden, 0.36 for the Netherlands and 0.33 for Finland, where all results are averages of estimated coefficients over age categories. We ran the same regressions presented above, with the addition of a gender dummy variable equal to one for the female descendant and zero for the male (results not shown). For the two age categories with descendants born before 1960, being female had a significant negative impact on the educational outcome of the descendant in the father regression. In the other models, this result holds only for the oldest age category (those born before 1945). The effect is negative but insignificant for all those born in or after 1960, a result documenting the steeper rise of the educational expansion for women. The inclusion of the gender dummy variable hardly influences the coefficients for the educational transmission of education; it changes the coefficients only at the third place after the decimal point. This result documents that the level of education and persistence of education over generations needs to be clearly distinguished in analyses of educational attainment. The fact that there is a strong educational expansion does not imply that mobility at the end of the expansion will be higher than before or that women are less or more dependent now on the education of their parents (father and/or mother), even in light of the fact that the educational gap is closing. We will therefore turn to our main goal, analyzing the relationships of same-gender and cross-gender intergenerational transmission of educational attainment. One well established approach to gain further evidence on intergenerational transmission is the Markovian approach¹¹.

¹¹For Markovian approach theory relevant to intergenerational transmissions and transfers, see e.g. Shorrocks (1978), Geweke (1986), and Van de Gaer (2001). See Norris (1997) for theory on Markov Chains.

1.3.3 Markovian Approach

In this section we calculate right stochastic matrices for the transitions of the Markovian process which describes the intergenerational educational transmission. For the reader's convenience we recall the basic framework as well as the basic measurement issues concerning the markovian approach for analysing intergenerational transmission of education.

Let \mathcal{E} be a finite state space, where $e_i \in \mathcal{E}$ are the states and e is the number of states. Let $P = [p_{ij}] \in \mathbb{R}_+^{e \times e}$ be a stochastic matrix where the probability of moving from state e_i to state e_j is defined as $Pr(j|i) = p_{ij} \geq 0$, given by the element in row i and column j of the matrix P . Of course, $\sum_{j=1}^e p_{ij} = 1$, meaning that every origin state leads to some final state with probability 1.

In the present case, the states e_i are given by the set of different educational levels. E^f denotes the row vector which gives the marginal distribution of the education levels of the fathers, E^m denotes that of the mothers, and E^d is that of the descendants. Therefore, a row vector $p_{i1}, p_{i2}, \dots, p_{ie}$ is the educational "lottery" faced by a descendant whose father or mother belongs to educational class i .

The survey question provides six different school levels¹², which we aggregated into three classes of educational attainment¹³, already defined in table 1.1. Given the transition matrix, we calculate several different measures of mobility. Consider the following simple example which illustrates the intuition for this approach:

¹²1. No degree; 2. Compulsory school level; 3. Apprenticeship or vocational school degree; 4. Medium-level or technical school; 5. Matura and higher level vocational school; 6. University, Fachhochschule

¹³We call the classifications "low" (1-3), "medium" (4-5) and "high" education (6). For a detailed discussion of the Austrian Educational System in an economic context, see Fersterer (2001).

Example. Let us suppose a simple example in which we have a population of six fathers and six descendants. Education levels are only low or high. Three fathers have low education and three fathers have high education, while three descendants have low education and three descendants have high education. Let us assume that one descendant has more education than her father and one descendant has less education than her father. The transition probability is given by $Pr(j|i) = p_{ij} = w_{ij} / \sum_{j=1}^e w_{ij}$, where w_{ij} is the sum of the weights for father-descendant pairs associated with educational transition from educational class i to class j for $i, j = 1, 2, \dots, e$. The associated transition matrix P is therefore given by

$$P = \begin{bmatrix} p_{1,1} & p_{1,2} \\ p_{2,1} & p_{2,2} \end{bmatrix} = \begin{bmatrix} 2/3 & 1/3 \\ 1/3 & 2/3 \end{bmatrix}$$

which gives the transition from the educational distribution of the fathers population to the educational distribution of the descendant population. In this case, it is

$$\underbrace{\begin{bmatrix} 3 & 3 \end{bmatrix}}_{E^f} \times \underbrace{\begin{bmatrix} 2/3 & 1/3 \\ 1/3 & 2/3 \end{bmatrix}}_P = \underbrace{\begin{bmatrix} 3 \\ 3 \end{bmatrix}}_{E^d}.$$

We use HSHW 2008 data to construct analogous vectors of educational distributions. The vectors E^f , E^m , and E^d and the corresponding transition matrices (by rows and columns) $P^{f \rightarrow d}$ (table 1.3) and $P^{m \rightarrow d}$ (table 1.4) are ordered from high (e_1) to low (e_3) education levels.

The transition matrix $P^{f \rightarrow d}$ shows that, for example, a descendant whose father holds the highest degree has a 0.5 probability of holding also the highest degree, and a 0.95 probability of obtaining at least a medium level education. For a descendant of a

Table 1.3: Transition Matrix of Fathers and their Descendants

		Descendant's Education		
		High	Medium	Low
Father's Education	High	0.50	0.45	0.05
	Medium	0.24	0.54	0.22
	Low	0.06	0.24	0.70

Table 1.4: Transition Matrix of Mothers and their Descendants

		Descendant's Education		
		High	Medium	Low
Mother's Education	High	0.45	0.50	0.05
	Medium	0.29	0.52	0.19
	Low	0.07	0.26	0.67

father with a "low" level of education, the same probabilities are 0.06 and 0.3, respectively. The transition matrix $P^{m \rightarrow d}$ appears to exhibit fairly similar trends. However, summing up the diagonal of the matrices gives 1.74 for $P^{f \rightarrow d}$ and 1.64 for $P^{m \rightarrow d}$, providing the first hint for stronger persistence of descendants' education coming from fathers' level of education.

One way of ordering the lotteries that any two descendants are facing given their parents' education is the stochastic dominance ordering. Let p_i denote the row vector of the i th row of a right stochastic transition matrix P . Let us assume an "at least as good as" preference relation \geq on educational lotteries. In the sense of stochastic dominance the lottery p_i is "at least as good as" lottery p_j if $p_{i,1} + p_{i,2} + \dots + p_{i,k} \geq p_{j,1} + p_{j,2} + \dots + p_{j,k} \forall k = 1, 2, \dots, e - 1$ and "better" ($>$) if at least one inequality holds. In the case of $P^{f \rightarrow d}$ (and $P^{m \rightarrow d}$) this translates to $p_1 > p_2 > p_3$. Therefore, the transition matrix is said to be monotone because $\forall i = 1, 2, \dots, e - 1, \sum_{j=1}^k p_{i,j} \geq \sum_{j=1}^k p_{i+1,j}, \forall k = 1, 2, \dots, e - 1$. Put simply, let us choose two people from the descendant population whose fathers have different education levels. The following statement is always true: The one with the more highly educated father faces a "better" lottery in the stochastic dominance sense.

To investigate the cross- and same-gender patterns of the transmission of educational attainment, we calculate the following transition matrices, all of which turned out to be monotone: $P^{f \rightarrow d_d}$, $P^{f \rightarrow d_s}$, $P^{m \rightarrow d_d}$ and $P^{m \rightarrow d_s}$ where d_d and d_s are the female (daughters) and male (sons) subsets, respectively, of the set of the descendants population.

Mobility Measures. Shorrocks (1978) provides a general framework for measuring mobility when data are provided in the form of a transition matrix. In general, those measures can be defined as continuous real functions of the form $M(\cdot) : \mathcal{P} \mapsto \mathbb{R}$ over the set of transition matrices \mathcal{P} .

Generally, there are two ways of analyzing mobility: mobility as *movement* and mobility as *independence*. In the former, a measure of mobility prefers those mobility matrices which incorporate more movement to those which incorporate less movement. If mobility is defined as independence, on the other hand, a mobility measure prefers those mobility matrices which incorporate less unequal lotteries to those which incorporate more unequal lotteries. In this sense, independence can also be interpreted as "equality of opportunity".

To follow an independence approach, which requires that the highest mobility is achieved if a matrix induces perfect origin independence, it is convenient to assert that $M(I) \leq M(P) \leq M(\bar{P})$, where $I \in \mathcal{P}$ is the identity matrix, $P \in \mathcal{P}$ is any transition matrix, and $\bar{P} \in \mathcal{P}$ is a transition matrix whose rows are identical. The identity matrix generates no transition between states and should be assigned with the least level of mobility. The matrix $\bar{P} \in \mathcal{P}$, on the other hand, should be assigned the highest level of mobility, because it induces perfect origin independence (Fields and Ok, 1996; Prais, 1955). Of course, this property is not always desirable- especially when mobility is

defined as movement. However, for an intergenerational framework, such a conception is relevant because we consider mobility to be independence. For convenience, the measures are normalized to the interval $[0, 1]$. Van de Gaer et al. (2001) show that because the axioms introduced by Shorrocks (1978) are inconsistent on the full domain of \mathcal{P}^{14} , the standard measures are not appropriate to measure mobility defined as independence on the full domain of \mathcal{P} . Van de Gaer et al. (2001) introduce suitable measures for the full domain of \mathcal{P} but since we only have to deal with monotone transition matrices, we can restrict our set to $\Xi \subset \mathcal{P}$, which contains only monotone transition matrices in order to be able to use conventional measures (Fields and Ok, 1996; Van de Gaer et al., 2001).

Pursuing an analysis of the independence measurement in our transition matrices, we turn to four related but unique tools. One widely used measure of the independence family of indices is the Second Eigenvalue Index. The eigenvalues of a given transition matrix ordered by the absolute value of their real part are given by $\lambda_i = |\lambda_1| \geq |\lambda_2| \geq \dots \geq |\lambda_n|$. Every transition matrix has $\lambda_1 = 1$. The Second Eigenvalue Index measures the distance of any given transition matrix to the origin independent matrix \bar{P} ; it is given by $M^{SE}(P) \equiv 1 - |\lambda_2|$. If λ_2 is equal to zero, then the transition matrix is equivalent to the limiting origin independent matrix. Therefore M^{SE} equals 1 when the outcome distribution is independent of the original distribution. If, on the other hand, M^{SE} equals 0, then the educational attainment of the descendant population is perfectly determined by the educational attainment of the parent population.

A second measure in this family of indices is the measure proposed by Shorrocks

¹⁴The relevant axioms are

- (i) *Monotonicity*: $P > P'$ when $p_{ij} \geq p'_{ij} \forall i \neq j$ and $p_{ij} > p'_{ij}$ for some $i \neq j$. Therefore $M(P) > M(P')$.
- (ii) *Immobility*: $M(I) = 0$. Minimum should be reached for identity matrix.
- (iii) *Perfect Mobility*: Let $P'' = (1/n)uu'$ where u is an n -dimensional vector of ones. Then $\forall P \neq P'' \in \mathcal{P}$ it follows that $M(P'') > M(P)$

Clearly (i) and (iii) are inconsistent on the domain of \mathcal{P}

(1978)¹⁵. Based on the trace of the transition matrix, this index evaluates the concentration around the diagonal of the matrix: $M^S(P) \equiv \frac{e - \text{trace } P}{e - 1}$. We use the Determinant Index, given as $M^D(P) \equiv 1 - |\det(P)|^{1/n-1}$ as our third index fulfilling Shorrocks axioms on Ξ . It is related to the average magnitude of the moduli of the eigenvalues of P .

The three indices above provide no indication of the number of classes an average descendant stands away from the educational class of his or her parent. Therefore we also take a look at an ad-hoc measure which does so. The so called absolute average jump $AAJ(P)$ gives the mean number of classes moved in absolute value. In our case, $AAJ(P) \in [0, 2]$. Table 1.5 shows the selected mobility indices¹⁶ for all described transition matrices.

Table 1.5: Mobility Measures for same-gender and cross-gender parent-descendant transition matrices

	Mobility Measures			
	Second Eigenvalue	Shorrocks	Determinant	Absolute Average Jump
Father's to Sons	0.402 (4)	0.588 (4)	0.632 (4)	0.441 (4)
Father's to Daughters	0.441 (2)	0.657 (2)	0.733 (2)	0.464 (2)
Mother's to Sons	0.463 (1)	0.765 (1)	0.811 (1)	0.564 (1)
Mother's to Daughters	0.437 (3)	0.626 (3)	0.677 (3)	0.453 (3)

Indices with rank in parentheses (1=most mobile)

All of the indices lead to exactly the same ranking, which is quite remarkable for these indices and is quite a strong result. The most mobile, and therefore the most independent, is the relationship between the education of sons in relation to that of their mothers. The next most independent is females in relation to their fathers. The

¹⁵Sometimes referred to as Shorrocks Mean Exit Time or Prais Index.

¹⁶Of course there exists a wide range of further indices within the family of indices which fulfill Shorrocks (1978) axioms on Ξ and also a number of further ad-hoc indices and other measures of rank-correlation. For better clarity we decided to use a very restricted set of them. Most of the indices we tried led to a equal or very similar ranking and all ranked identical over the set $\{P^{f \rightarrow d_s}, P^{f \rightarrow d_d}, P^{m \rightarrow d_s}, P^{m \rightarrow d_d}\}$.

strongest dependence is between male descendants and their fathers. Clearly, males are more dependent on their fathers than on their mothers and females are more dependent on their mothers than on their fathers¹⁷. There is a strong gender-specific pattern in the transmission of educational attainment: children are likely to follow in the footsteps of their same-gender parent.

This finding has significant implications when thinking about the persistence of gender inequality. Not only is there a strong intergenerational persistence in educational outcomes that reinforces divisions across families and classes in society, but those divisions are further deepened on the level of each individual family along gendered lines. We might consider the intergenerational persistence of educational attainment to be on a *macro* level of inequality, while the gendered persistence within families might be seen as a *micro* level effect.

1.3.4 Multivariate Analysis - Ordered Logit

In order to obtain further evidence on the gender issue as well as to check for robustness of results, we conduct a multivariate ordered logit estimation¹⁸, similar to the method of Bauer and Riphon (2004) and Daouli et al. (2008). Educational attainment of the descendant (E^d) split by gender are the dependent variables, while educational attainment of the fathers (E^f) and mothers (E^m) along with the age of the descendant are the independent variables. For the sake of including as many observations as possible, we pool the ages of the descendant instead of estimating the different descendant population subsets¹⁹. The modes of the educational attainment are excluded for fathers'

¹⁷Of course in any approach where the educational outcome of a descendant only depends on either his or her mother or father, we can't control for assortative mating. This is done in the ordered logit approach in section 1.3.4

¹⁸The possible use of a multinomial logit or generalized ordered logit approach which could be favored leads to similarly interpretable results.

¹⁹If one uses the subsets for estimation transmissional coefficients are, as expected, lower for the younger subsets.

and mothers' education to serve as reference category. For both mothers and fathers, the mode is e_3 , the lowest educational category in our classification. Table 1.6 shows the marginal effects of the ordered logit estimation evaluated at the means and modes.

The probabilities for an average descendant to obtain educational levels low, medium and high are given in the first row of table 1.6. Having a father with a high education level instead of low education level increases *ceteris paribus* the probability (in absolute terms) of holding a high educational degree by 0.406 (40.6 percentage points; significant at 1% level) for sons and 0.261 (26.1 percentage points; significant at 1% level) for daughters. Having a mother with a high education level instead of a low education level has insignificant effects for sons, and increases a daughter's probability of holding a medium or high educational degree by 0.246 (24.6 percentage points; significant at 1% level) and 0.087 (8.7 percentage points; significant at 10% level), respectively. All the significant marginal effects have the expected signs: descendants with more highly educated parents are more likely to have high levels of education themselves. Being older leads to a lower probability of higher education and higher probability of lower education, but these effects are significant only for females, a result that documents the higher importance of the educational expansion for females and may be interpreted as a hint for other important determinants (e.g. the feminist movement, reforms, etc.) for female education in this period for which there are no controls in the model.

A father's level of education has a stronger effect than a mother's level of education on descendants of both genders. Overall, the effect of mothers' education levels on the educational attainment of daughters is substantially higher than the effect for their sons. Overall, the effect of fathers' education on their sons' education is higher than that on their daughters. Effects are stronger when the parent population has a high, rather than medium, level of education.

Table 1.6: Marginal Effects at Means (Modes) for Ordered Logit Estimations on Female and Male Subsets

	<i>descendant_{low}</i>		<i>descendant_{medium}</i>		<i>descendant_{high}</i>	
	male	female	male	female	male	female
$Pr(Y X)$	0.714	0.703	0.230	0.265	0.057	0.033
<i>father_{medium}</i>	-0.374 (0.046)***	-0.356 (0.042)***	0.206 (0.022)***	0.258 (0.028)***	0.169 (0.034)***	0.098 (0.021)***
<i>father_{high}</i>	-0.566 (0.045)***	-0.542 (0.043)***	0.160 (0.045)***	0.281 (0.031)***	0.406 (0.082)***	0.261 (0.061)***
<i>mother_{medium}</i>	-0.157 (0.054)***	-0.232 (0.051)***	0.107 (0.034)***	0.182 (0.037)***	0.050 (0.021)**	0.050 (0.015)***
<i>mother_{high}</i>	-0.177 (0.132)	-0.333 (0.110)***	0.119 (0.079)	0.246 (0.062)***	0.058 (0.054)	0.087 (0.050)*
<i>descendant_{age}</i>	0.001 (0.001)	0.004 (0.001)***	-0.001 (0.001)	-0.003 (0.001)***	-0.000 (0.000)	-0.001 (0.000)***

Notes:

- (i) Males Regression: No.Obs.=900; Log Likelihood=-730.008; Cox-Snell R2=0.22; Nagelkerke R2=0.26; McFadden R2=0.12
- (ii) Females Regression: No.Obs.=992; Log Likelihood=-763.382; Cox-Snell R2=0.25; Nagelkerke R2=0.29; McFadden R2=0.15
- (iii) *, **, *** denotes significance at 10%, 5%, 1% level; Standard errors are given in parentheses

To illustrate some of these results, figures 1.2 to 1.7 show the probabilities of reaching low, medium and high education for 6 benchmark cases of different combinations of parental educational attainment over the age of the descendants calculated from the ordered logit models presented in table 1.6. Comparing figures 1.4 to 1.5 and 1.6 to 1.7, we see that a jump in the fathers' education from low to high leads to a higher probability of having a high education level for sons (in relation to daughters) and that a change in the mothers' education level from low to high leads to higher chances of attaining medium or high levels of education for daughters (in relation to sons). The result of the stronger same-gender than cross-gender relationships holds if we conduct the regressions on two different subsets of the sample according to the age-category of the descendants referring to the time before the educational expansion (age categories born before 1945 and between 1945 and 1959) and during (and after) the educational expansion (age categories born between 1960-1969 and between 1970-1983), with the sole exception of father-son vs. father-daughter pairs for the older age categories subset. That means that even as the differences in educational outcomes in relation to gender are diminishing over time (see section 1.3.1 and 1.3.2), which is documented in many studies (see e.g. Steiner 1998; Biffl 2002; Bacher 2003; Spielauer 2004), same-gender versus cross-gender differences in the intergenerational transmission of educational attainment seem to be quite stable to date. Note that we can neither show the influence of fathers on their daughters versus the influence on their sons nor the influence of mothers on their sons versus the influence on their daughters to be statistically significantly different at the 5% level in a nested model, which may also have to do with different distributional patterns over time - as they were discussed in section 3.2 - and of course due to assortative mating and the limited number of observations. Nevertheless, the robustness of the results in different approaches, i.e. in the Markovian approach as

well as the consistently different patterns in the ordered logit approach over different sample subsets (female descendants, male descendants over different age categories) can be viewed as strong evidence for the differences in same-gender versus cross-gender relationships of intergenerational educational attainment.

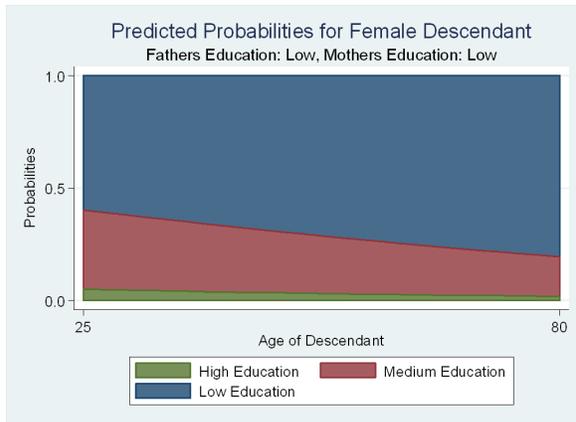


Figure 1.2: FEMALES; Mother: Low; Father: Low

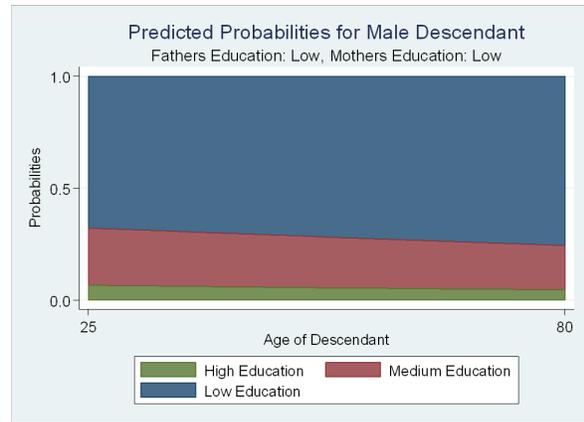


Figure 1.3: MALES; Mother: Low; Father: Low

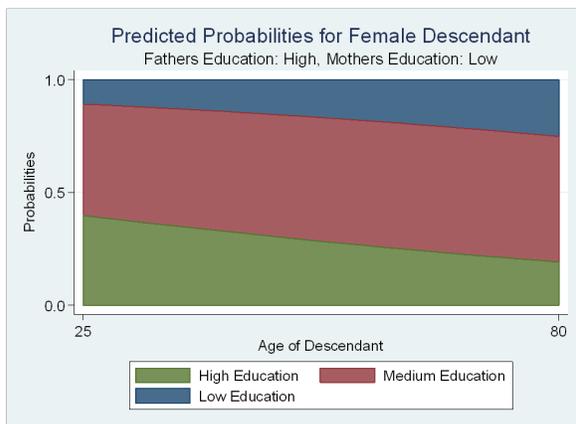


Figure 1.4: FEMALES; Mother: Low; Father: High

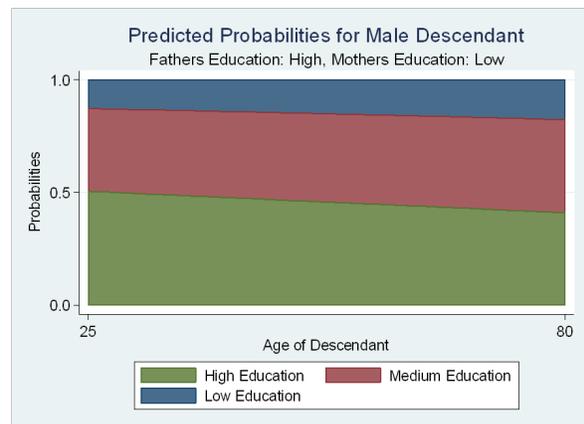


Figure 1.5: MALES; Mother: Low; Father: High

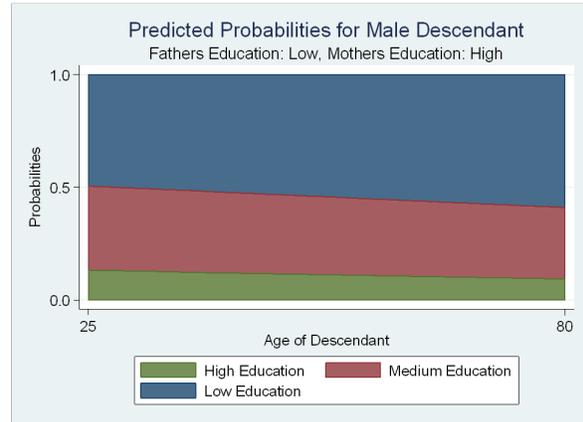
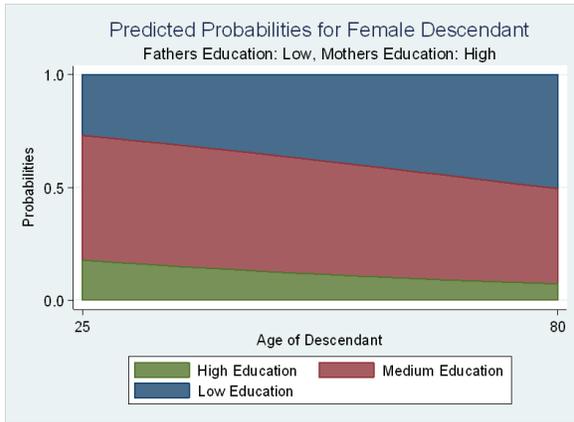


Figure 1.6: FEMALES; Mother: High; Father: Low
Figure 1.7: MALES; Mother: High; Father: Low

1.4 Conclusions

As far as the authors know the literature, this is the first paper to focus specifically on the gendered patterns of the intergenerational transmission of educational attainment. We confirm the results of previous studies, documenting a persistence in the intergenerational transmission of educational attainment (see Hertz et al. 2007 for International Results and Spielauer 2004 for some results on Austria). We tested the intergenerational persistence in several ways, all of which gave robust results. Therefore, we can confidently assert that there is a connection between the educational level of one's parents' and one's own educational achievements. These results suggest that a more equitable society will provide resources for children from less fortunate families so that they may catch up with their peers educationally, because those children face barriers to educational attainment without such public support. We assert that the state should serve as a bridge for children from lower-income and lower-educational backgrounds to educational achievement, when their parents cannot do so, and reference the educational reforms in Austria as an example of the positive effects of publicly supported education.

We also confirm the gap between male and female educational attainments, which has been diminishing over the past century - especially since the educational expansion beginning in the 1960s. We think that the general educational expansion and the diminished gender-gap are due, at least in part, to educational reforms that expanded access to the classroom. The fact that we see a decline in the educational gap starting soon after the expansion in publicly funded education and continuing over several decades suggests that eliminating or at least lowering private costs to acquiring human capital can help reduce the inequality in educational attainment. We also credit the feminist movement, which has pushed for women and girls' access to schooling and general equality with males, while simultaneously encouraging people to question and rethink traditional gender roles which had dictated to women that they would not benefit from obtaining an education.

Our main contribution to the literature and field of feminist economics is the exposure of the gendered differences in the process of intergenerational transmission of educational attainment. To supplement other studies that lack this important feature (e.g. Spielauer 2004), we analyze the differences between the influence of fathers versus mothers educational attainment on educational attainment of descendants, finding the result that the education of the father is more influential for both sons and daughters, in absolute terms. Still, the effect of mothers' education on their daughters' education is substantially higher than its effect on the sons' education, and the effect of fathers' education on the sons' education is stronger than on their daughters' education. The gendered result is also consistent over the different approaches we used, namely the Markovian as well as the ordered logit approach, the latter allowing us to control for assortative mating. To this end, we believe that the intergenerational transmission of gender roles might be playing a significant part in the outcome of children's education. This topic is one for further research and discussion.

To the extent that the results are comparable, the level of intergenerational educational attainment correlation seems to be higher in Austria than in Northern European countries such as The Netherlands, Finland, and Sweden and closer to Southern European countries such as Italy or Slovenia. The division between secondary modern school (Hauptschule) and grammar school (Allgemeine Höhere Schule) at the age of 10 might be relevant for a relatively low level of educational mobility (see Spielauer 2004) up to date. Furthermore the low level of women in higher positions at the upper end of the educational system - university assistants and professors - is very worrisome.

The intergenerational persistence of an outcome such as education stands as a serious threat to the equity of a society. The results that this persistence (i) varies for different countries, (ii) varies over time, and (iii) is constructed along gender lines, serve as strong evidence in favor of the argument that the persistence is clearly a social phenomenon. We are encouraged by the finding that the educational gender gap as well as the persistence of intergenerational educational transmission has been diminishing over time, and recognize the importance of strong legislation and political activism in promoting its continual decline. Concerning the gendered aspects of intergenerational educational transmission, more research utilizing far bigger datasets is necessary to come to conclusions about how strong it is and how it has developed over time. At the least, we found strong evidence of its existence and importance for educational outcomes of the existing Austrian population and some evidence that it is still relevant for the younger cohorts.

1.5 Literature

Bacher, Johann. 2003. Soziale Ungleichheit und Bildungspartizipation im weiterführenden Schulsystem Österreichs. *Österreichische Zeitschrift für Soziologie* 28(3): 3-32.

Bauer, Philipp and Regina T. Riphon. 2004. Heterogeneity in the Intergenerational Transmission of Educational Attainment: Evidence from Switzerland on Natives and Second Generation Immigrants IZA Discussion Paper No. 1354.

Belzil, Christain and Jörgen Hansen. 2003. Structural Estimates of the Intergenerational Educational Correlation. *Journal of Applied Econometrics* 18(6): 679-696.

Becker, Gary S. and Tomes, Nigel. 1979. An Equilibrium Theory of the Distribution of Income and Intergenerational Mobility. *The Journal of Political Economy* 87(6): 1153-1189.

—. 1986. Human Capital and the Rise and Fall of Families. *Journal of Labor Economics* 4(2), part 2: S1-S39

Biffi, Gudrun. 2002. Der Bildungswandel in Österreich in den Neunziger Jahren. *Wifo-Monatsberichte* 6/2002: 377-384.

Behrman, Jere R., and Mark R. Rosenzweig 2002. Does Increasing Women's Schooling Raise the Schooling of the Next Generation? *American Economic Review*, 92(1): 323-334.

Black, Sandra E., Paul J. Devereux, and Kjell G. Salvanes. 2005. Why the Apple Doesn't Fall Far: Understanding Intergenerational Transmission of Human Capital. *American Economic Review* 95(1): 437-449.

BMWF (Bundesministerium für Wissenschaft und Forschung). 2007. Statistisches Taschenbuch. Bundesministerium für Wissenschaft und Forschung.

Bourdieu, Pierre. 1984. Die feinen Unterschiede: Kritik der gesellschaftlichen Urteilskraft. Suhrkamp Taschenbuch Wissenschaft, Frankfurt.

Bowles, Samuel and Herbert Gintis. 2001. The Inheritance of Economic Status: Education, Class, and Genetics. Working Paper 01-01-005, Santa Fe Institute.

Bowles, Samuel, Herbert Gintis, and Melissa Osborne Groves [eds]. 2005. Unequal Chances: Family Background and Economic Success. New Jersey: Princeton University Press.

Checchi, Daniele, Carlo V. Fiorio, and Marco Leonardi. 2008. Intergenerational Persistence in Educational Attainment in Italy. IZA Discussion Paper No. 3622.

Chevalier, Arnaud, Kevin Denny, and Dorren McMahon. 2003. A Multi-Country Study of Intergenerational Educational Mobility. School Of Economics, University College Dublin, Working Paper No. 200314.

D'Addio, Anna. 2007. Intergenerational Transmission of Disadvantage: Mobility or Immobility Across Generations? OECD Social Employment and Migration Working Paper No. 52.

Dardanoni, Valentino, Antonio Forcina, and Salvatore Modica. 2008. Direct Causal Effects in Education Transmission. submitted, available at <http://dse.univr.it/it/documents/it4/dfm-final.pdf> (accessed May 2009).

Daouli, Joan, Michael Demoussis, and Nicholas Giannakopoulos. 2008. Mothers, Fathers and Daughters: Intergenerational Transmission of Education in Greece. Available at SSRN: <http://ssrn.com/abstract=1161014>. (accessed June 2009).

Dearden, Lorraine, Stephen Machin, and Howard Reed. 1997. Intergenerational Mobility in Britain. *The Economic Journal* 107(440): 47-66.

Fersterer, Josef. 2001. Erträge der Schulbildung in Österreich - Eine empirische Untersuchung. Dissertation, University of Linz, Austria.

Fessler, Pirmin, Peter Mooslechner, and Martin Schürz. 2012. Intergenerational Transmission of Educational Attainment in Austria, *Empirica*: Volume 39, Issue 1 (2012), Page 65-86.

Fields, Gary S. and Efe A. Ok. 1996. The Measurement of Income Mobility: An Introduction to the Literature. Working Paper 96-05, C.V. Starr Center for Applied Economics, New York University.

Galton, Francis. 1877. Typical Laws of Heredity. *Nature* 15: 512-4.

Geweke, John, Robert Marshall, and Gary A. Zarkin. 1986. Mobility Indices in Continuous Time Markov Chains. *Econometrica* 54(6): 1407-1423.

Hansen, Karsten T., James J. Heckman, and K. J. Kathleen Mullen. 2004. The Effect of Schooling and Ability on Achievement Test Scores. *Journal of Econometrics* 121(1-2): 39-98.

Haveman, Robert and Barbara Wolfe. 1995. The Determinants of Children's Attain-

ments: A Review of Methods and Findings. *Journal of Economic Literature* 33(4): 1829-1878.

Heckman, James J. 2008. Schools, Skills, and Synapses. IZA Discussion Paper No. 3515.

Herrnstein, Richard J. and Charles Murray. 1994. *The Bell Curve*. New York: Free Press.

Hertz, Tom; Tamara Jayasundera, Patrizio Piraino, Sibel Selcuk, Nicole Smith, and Alina Verashchagina. (2007). The Inheritance of Educational Inequality: International Comparisons and Fifty-Year Trends, *The B.E. Journal of Economic Analysis and Policy* 7(2): Article 10.

Holmlund, Helena, Mikael Lindahl, and Erik Plug. 2008. The Causal Effect of Parent's Schooling on Children's Schooling: A Comparison of Estimation Methods. IZA Discussion Paper No. 3630.

Landler, Frank. 1997. *Das österreichische Bildungswesen in Zahlen - Analyse und Computersimulation des Schulsystems und der Qualifikationsstruktur der Bevölkerung*. Vienna, WUV-Universitätsverlag.

Loehlin, John C. 2005. Resemblance in Personality and Attitudes Between Parents and Their Children: Genetic and Environmental Contributions. In *Unequal Chances: Family Background and Economic Success*. Bowles, Samuel, Herbert Gintis, and Melissa Osborne Groves [eds]. New Jersey: Princeton University Press.

Moen, Phyllis, Mary Ann Erickson, and Donna Dempster-McClain. 1997. Their Mother's Daughters? The Intergenerational Transmission of Gender Attitudes in a World of Changing Roles. *Journal of Marriage and the Family* 59(2): 281-93.

Mulligan, Casey B. 1999. Galton versus the Human Capital Approach to Inheritance. *The Journal of Political Economy* 107(6) part 2: S184-S224.

Norris, James R. 1997. *Markov Chains*. Cambridge, England: Cambridge University Press.

OECD. 2008. *Growing Unequal*. Chapter 8: Intergenerational Mobility, pp. 203-221.

Osborne Groves, Melissa A. 2005. Personality and the Intergenerational Transmission of Earnings from Fathers to Sons. In *Unequal Chances: Family Background and Economic Success*. Bowles, Samuel, Herbert Gintis, and Melissa Osborne Groves [eds]. New Jersey: Princeton University Press.

Plug, Erik. 2004. Estimating the Effect of Mother's Schooling Using a Sample of Adoptees. *The American Economic Review* 94(1): 358-368.

Prais, Sig J. 1955. Measuring Social Mobility. *Journal of the Royal Statistical Society Series A*, 118(1): 56-66.

Shorrocks, Anthony F. 1978. The Measurement of Mobility. *Econometrica* 46(5): 1013-1024.

Spielauer, Martin. 2004. Intergenerational Educational Transmission within Families: An Analysis and Microsimulation Projection for Austria. *Vienna Yearbook of Population Research*.

Steiner, Mario. 1998. Empirische Befunde zur Chancengleichheit im Österreichischen Bildungssystem. *Erziehung Heute*, 4/98, Innsbruck.

Van de Gaer, Dirk, Erik Schokkaert, and Michael Martinez. 2001. Three Meanings of Intergenerational Mobility. *Economica* 68(272): 519-537.

Wagner, Karin and Siegfried Zottel. 2009. OeNB-Immobilienvermögenserhebung 2008 der privaten Haushalte - Beschreibung, verwendete Methoden und Bewertungsansätze. in *Statistiken Q4/09*: 45-65.

Chapter 2

Assets and Liabilities of Austrian Households

2.1 Introduction

The liabilities of private Households have risen sharply in many OECD countries since the mid eighties (Girouard et al. 2006). At first sight, one would think that debt-holding households (*i*) hold no or very little financial wealth and/or (*ii*) are low income households. In other words, one might understand debt to come about as a substitute in the face of a lack of financial resources. In this case, we are presented with a certain danger to financial stability: households holding much debt would be quite vulnerable to adverse financial shocks, without any financial wealth to absorb such shocks. A macro level analysis shows that both liabilities and assets are rising, but it is unclear how this rise translates to the micro (household) level.¹

Macroeconomic analyses only allow us to look at the expansion of both on an aggregate level. There are two main problems with a macro-level analysis of financial assets (henceforth referred to as assets) and liabilities: (*i*) it is rather difficult to separate true private households

¹Some descriptive parts of this study are already published (Fessler and Mooslechner 2008).

from Non-Governmental Organizations, churches and other such organizations in the National Accounts, and (ii) we cannot know which households hold debt and which households hold financial wealth holdings - nor can we know to what extent some of the households hold both debt and financial wealth. We therefore need to analyse debt and asset holdings on the micro level to analyse threats of debt to financial stability; put simply, if households holding debt also hold certain amounts of financial assets, the risk exposure towards financial shocks is quite low. If debt and asset holders are separate groups, the risks the risks related to financial shocks are quite high.

Cox et al. (2002) use a descriptive analysis to show that in Great Britain, the relation between liabilities and assets in private households decreases as asset values increase, but in absolute terms, households with the highest wealth holdings also hold the highest amount of debt. Aizcorbe et al. (2003) present similar results for the USA (see also Barnes and Young (2003), Barwell et al. (2006) for the UK or Herrala (2006) for Sweden), where participation on the credit markets and the amount of debt held by a household rises with income and wealth. Beer und Schürz (2007) conducted a similar study for Austria, finding that higher income household have more participation in the credit market, although their analysis of the effect of wealth holding on debt are unclear.

It is not completely clear why households would hold large amounts of debt and large amounts of financial wealth at the same time. In general, holding debt is expensive, because of the interest payments due on the debt. The only situation in which it would be rational to hold both debt and assets would be when the returns to assets are higher than the cost of debt. If this is not the case, the rational household should use the assets to pay off the (more expensive) debt. A situation in which returns to assets are higher than interest payments on debt is, of course, not sustainable in the long run, so we must look for other explanations. One possibility is that households have a preference for liquidity, and therefore accept the cost of holding debt in order to have the liquidity provided by the assets. Another explanation is

that their asset holdings are too illiquid for use in cutting down debt.

Recent behavioral economics literature analyzing data on credit card debt and financial wealth holdings of US households offers another explanation for the phenomenon of holding both assets and expensive debt simultaneously (s. Laibson et al. 2000, Laibson 1997). Laibson et al. 2000 suggest that households face hyperbolic discounting, meaning that the discount rate for the near future is higher than the long term discount rate (dynamically inconsistent time preferences), and that encourages households to delay eliminating their debt.

This paper analyzes the extent to which households hold both assets and liabilities. It is organized as follows. Section 2.2 provides an overview of the correlation between assets and liabilities on macro level. In Section 2.3 we present the socioeconomic characteristics of households holding consumer and housing debt. While we present some international data in section 2.4 we concentrate on Austria in section 2.5 and employ a Heckman Selection model to investigate the determinates of holding liabilities, both probability and level. Section 2.6 concludes.

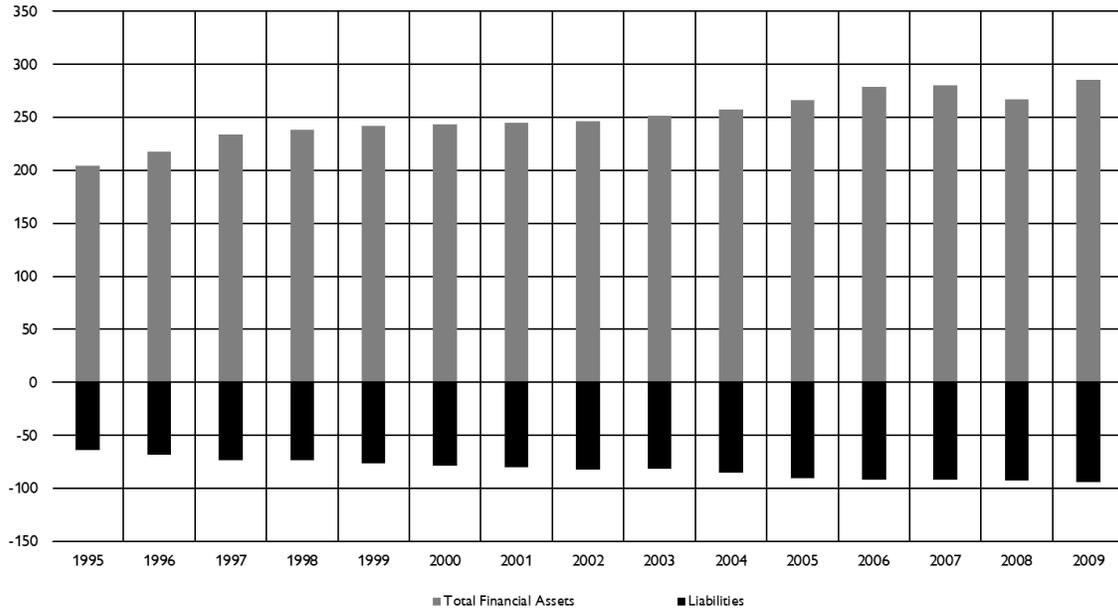
2.2 The macro level

All financial transactions have two sides. Financial Assets on one side imply liabilities on the other side. In most cases, the household sector has a structural surplus, meaning that they have higher financial assets than liabilities, and the firm sector has a structural deficit, meaning that they have higher liabilities than financial assets. In a closed economy, financial assets and liabilities cancel each other out over all sectors. In an open economy, a deficit in one country implies a surplus in another country, but financial assets and liabilities over all countries sum to zero. It is therefore not surprising that both financial assets and liabilities can increase together.

Figure 2.1 shows the increase in both household financial assets and liabilities over the

Assets an Liabilities of Austrian Households

in % of disposable income



Source: Statistik Austria; OeNB.

Figure 2.1: Household financial assets and liabilities

past decade.

Macro data can not tell us anything about the distribution of assets and liabilities. A first guess would be that assets and liabilities function as substitutes; households with low income and low assets might need liabilities to finance their consumption. On the other hand, households often need loans to finance assets, such as real estate wealth, which would mean that a household might have both high real assets and liabilities at the same time. Nevertheless, economic theory suggests that large holdings of financial assets combined with large holdings of liabilities should be less common, because interest paid on debt should on average be higher than interest earned from financial assets (given the assumption that banks have some costs as well).

In this following analysis, we use microdata to assess the connection between financial assets and liabilities.

2.3 Debtors in Austria

In this section we examine the socioeconomic characteristics of indebted Austrian Households. We use the Survey of Austrian Households Financial Wealth (SHFW 2004), which is a representative dataset on the financial wealth of Austrian households². We distinguish between debt used to buy real estate (real estate debt) from debt used to consume other goods (consumer debt). The socioeconomic characteristics used are those of the so-called household head³.

The life-cycle hypothesis tells us that an important determinant of the demand for debt is age. Young individuals with a high marginal utility of consumption and with an expectation of rising income should have a high demand for debt. From a certain age onwards, one would expect that the demand for debt decreases. Financial wealth is also an important determinant for the demand for debt: an increase in one's budget for financial wealth should decrease the demand for debt. The more liquid the financial wealth is, the stronger this effect should be. However, this is not always the case; there are situations when financial wealth and debt go hand in hand. For example, a person interested in obtaining a loan (e.g. for buying real estate wealth) may need to show a certain level of wealth in order to do so. This situation - of simultaneous financial assets and liabilities (at least for a short time) - occurs in the presence of credit restrictions. We therefore see that the relationship of financial wealth and debt holdings is not unambiguously clear. Nevertheless we would in general expect a negative relationship as debt holders should reduce their debt over time and then (or already during that period) build up financial wealth again.

Most households hold debt in order to buy real estate wealth. Real estate debt, which most often take the form of mortgages, is generally higher than consumer debt. Most house-

²The dataset consists of 2556 observations and was sampled using a stratified clustered multistage random sample. Listwise deletion was used, i.e. controlling for non-random item non-response is therefore not feasible. For details see Beer et al. 2006

³In the survey, the term 'household head' refers to that individual in a household who is either identified as the household head by the interviewees or the individual 'who best knows about the financial situation of the household'

holds (household heads) buy their real estate wealth between the ages of 30 and 39. From this time onwards, we would expect them to start to repay their debt and to build financial wealth again. While it is possible that some households prefer to stay liquid and hold debt and financial wealth simultaneously, we generally expect a negative correlation between assets and liabilities for those holding liabilities. Debt should decrease over time while financial wealth should increase over time.

With respect to income, we expect that one's expectations of his/her future stream of income to be the most important factor in determining debt holdings. The steeper one's income expectations are, the higher the demand for debt should be. If the expectation about the future stream of income is flat, there is less need to smooth consumption over the life-cycle and the incentive to hold debt should be lower. Along with the expectations about the shape of the future stream of income, uncertainty related to the realization of future income could play an important role in determining debt holdings. If the probability of becoming unemployed in the future is relatively high, the demand for debt should be lower. It remains unclear how current household income should be related to the demand of debt. On the one hand, a low income could be related to expectations of a steep incline of the income in the future. On the other hand, the degree of uncertainty is typically higher for low income jobs than it is for high income jobs. Nevertheless, higher income should generally lead to the possibility of holding more debt.

As a proxy for an expectation of the future stream of income, education could be very important. Higher education corresponds with a greater probability of having a steeper incline of future income. Furthermore, access to credit is easier for highly educated people because barriers to entry are smaller. Therefore, we expect education to be positively related to household debt.

Regional differences could also play an important role for the holdings of debt of households. If lenders are less developed in certain areas, this may induce lower demand for debt because

of higher entry barriers. Further, different regions could have different rates of unemployment, which might reflect the probability of becoming unemployed and therefore uncertainty concerning the future stream of income.

Banks verify the education level, income and wealth of their clients before they give them a loan. Therefore, the same socioeconomic characteristics influence the demand as well as the supply side of loans. Some households do not attempt to take out a loan, anticipating that they will be denied on the basis of their socioeconomic characteristics.

It seems likely that various other variables, which are difficult to observe, play an important role in the process of demanding and obtaining a loan. There may be historical and institutional reasons which could influence the attitude towards real estate wealth or debt as a whole.

2.3.1 Real estate debt

Approximately 29% of Austrian Households hold real estate debt. Holding debt to finance a house or an apartment is a more widely spread phenomenon than holding debt for other consumption issues; only about 17% of Austrian households hold consumer debt. The average real estate debt reaches a value of 57,708 Euro (Median: 36,000 Euro). See table 2.1 for an overview over the distributions. The ratio of households holding real estate debt is the highest in the 30 to 39 years age group. The probability of holding real estate debt rises with income. Regarding the share of real estate debtors over gross financial wealth deciles⁴, higher financial wealth corresponds to a higher share of debtors with a sharp decrease at the very top of the wealth distribution. It is only in the 10th decile that the share lies below the average share (table 2.2). Of course the share of households holding a real estate loan does not allow us to

⁴Gross financial wealth is defined as the sum of all values on checking accounts, savings accounts, all other saving vehicles, bonds, stocks, mutual funds, business assets and all payments to life insurance contracts. The inclusion of life insurance contracts does not lead to significant differences in all results of this paper.

Table 2.1: Percentile values of the reference group of debtors

Percentiles	Consumer debt (consumer debtors)	Real estate debt (real estate debtors)	All debt (all debtors)
10	1160	7000	5000
25	5000	15000	10000
Median	10000	36000	25000
75	20000	80000	50000
90	34000	140000	125000

conclude that there is a positive correlation between financial assets and the value of the debt. A first hint for such a relationship is given in table 2.3, which shows the average real estate debt and the average gross financial wealth for gross financial wealth deciles. It also shows that in higher wealth deciles, the values of real estate debts are higher, even when the average household could eliminate its debt given the amount of financial wealth it holds. In the first wealth decile debt is on average 18 times higher than the financial wealth of the households in that decile. On the other hand, the value of debt is just 1/5 of the average value of financial wealth in the highest decile. The burden for the households in the lower wealth deciles is obviously much higher than the burden for those in upper wealth deciles. This pattern could reflect a life-cycle effect. From a theoretical point of view, though, it remains unclear why households which could easily repay their debt do not do so, given the fact that they need to pay interest for their debt. If life-cycle theory is applicable, households with high values of financial gross wealth should not hold debt. Obviously, that is not the case, and it is therefore impossible to look at the debt of a household and conclude anything about its financial assets based on the level of debt. It is also impossible to identify rich or poor households in this way. The households with the highest debt are the households with the highest values of financial assets.

Of course it could be that bullet loans are the reason for such a picture. However, given the share of bullet loans in Austria, this explanation is by far not sufficient.

Table 2.2: Socioeconomic characteristics and real estate debt

		Share of real estate debtors	Difference to total value in percentage points
Age of the household head	18 - 29	29%	0
	30 - 39	42%	+13
	40 - 49	39%	+10
	50 - 59	34%	+5
	60 - 69	15%	-15
	70 - 79	6%	-24
	80+	15%	-14
Net income of household	up to EUR 749,-	9%	-20
	EUR 750,- to EUR 1.349,-	16%	-13
	EUR 1.350,- to EUR 2.249,-	27%	-2
	EUR 2.250,- to EUR 2.999,-	40%	+11
	EUR 3.000,- and more	40%	+11
Education of household head	Max. compulsory school	18%	-11
	Vocational school, middle school	30%	+1
	High school, grammar school	31%	+2
	University and equivalent	33%	+4
Household size	1 person	19%	-10
	2 persons	21%	-8
	3 persons	33%	+4
	4 persons	48%	+19
	5 and more persons	57%	+28
Gross wealth deciles	Decile 1	16%	-13
	Decile 2	28%	-1
	Decile 3	25%	-4
	Decile 4	27%	-2
	Decile 5	31%	+2
	Decile 6	34%	+5
	Decile 7	39%	+10
	Decile 8	36%	+7
	Decile 9	33%	+4
	Decile 10	22%	-7
occupation	Self-Employed (Freie Berufe)	29%	-10
	Entrepreneurs	32%	-7
	White-collar Workers	40%	+1
	Civil servants	39%	0
	Farmers	29%	-10
	Blue-collar Workers	41%	+2
Employed Total		39%	0
Total		29%	0

Table 2.3: Averages for real estate debt and gross financial wealth of real estate debtors

Gross financial wealth	Average real estate debt	Average gross financial wealth
Decile 1	28035	1537
Decile 2	54395	5334
Decile 3	45548	8980
Decile 4	41561	13853
Decile 5	55138	20057
Decile 6	60140	28084
Decile 7	69347	39885
Decile 8	59666	55468
Decile 9	63912	82849
Decile 10	83307	368456

2.3.2 Consumer debt

As mentioned above, approximately 17% of Austrian Households hold consumer debt. The average consumer debt amounts to 16,845 Euro (Median: 10,000 Euro) and is therefore lower than average real estate debt. Households holding consumer debt are on average younger than their real estate debt counterparts. It is also the case that households with consumer debt are often higher income households. Households with consumer debt are more likely to be found in the lower gross financial wealth deciles (table 2.4). However, even in the highest gross financial wealth decile, more than one out of ten households holds consumer debt. Looking at the average consumer debt versus the average gross financial wealth, the picture seems to be close to, but less pronounced than, that for real estate debt (table 2.5). The average gross financial wealth in the top decile is remarkably lower than for households with real estate debt. In general consumer debt seems to be less important for the upper part of the wealth distribution.

Credit constraints Looking at net financial wealth versus gross financial wealth, it is clear that the lowest net financial wealth decile holds higher gross financial wealth than the second lowest net financial wealth decile. This fact suggests that households in the lower part of the gross financial wealth distribution do not hold much real-estate debt. This fact could

Table 2.4: Socioeconomic characteristics and consumer debt

		Share of households holding consumer debt	Difference to total in percentage points
Age of household head	18 - 29	16%	-1
	30 - 39	27%	+10
	40 - 49	22%	+5
	50 - 59	18%	+1
	60 - 69	9%	-8
	70 - 79	3%	-14
	80+	5%	-12
Net income of household	up to EUR 749,-	7%	-10
	EUR 750,- to EUR 1.349,-	15%	-2
	EUR 1.350,- to EUR 2.249,-	16%	-1
	EUR 2.250,- to EUR 2.999,-	16%	-1
	EUR 3.000,- and more	24%	+7
Education of household head	Max. compulsory school	14%	-4
	Vocational school, middle school	18%	+1
	High school, grammar school	19%	+1
	University and equivalent	16%	-2
Household size	1 person	13%	-4
	2 persons	15%	-2
	3 persons	25%	+8
	4 persons	20%	+3
	5 and more persons	19%	+2
Gross financial wealth decils	Decile 1	26%	+9
	Decile 2	20%	+3
	Decile 3	22%	+5
	Decile 4	16%	-1
	Decile 5	14%	-3
	Decile 6	18%	+1
	Decile 7	10%	-7
	Decile 8	19%	+1
	Decile 9	15%	-3
	Decile 10	11%	-7
occupation	Self-Employed (Freie Berufe)	30.1%	+4.9
	Entrepreneurs	22.7%	-2.5
	White-collar Workers	22.8%	-2.4
	Civil Servants	22.7%	-2.5
	Farmers	6.0%	-19.2
	Blue-collar Workers	24.1%	-1.1
Employed total		23.1%	0
Total		17%	0

Table 2.5: Average consumer debt and average gross financial wealth of consumer debtors

Gross financial wealth decil	Average consumer debt	Average gross financial wealth
Decile 1	8130	1442
Decile 2	12390	5034
Decile 3	10502	9195
Decile 4	16969	14118
Decile 5	14865	20024
Decile 6	23969	28324
Decile 7	13536	38826
Decile 8	17854	55717
Decile 9	25446	81343
Decile 10	40094	282062

be explained by age effects or credit constraints (or both), because those households may have not enough gross wealth to get a mortgage. To take this fact into account in the econometric analysis, it would be necessary to have some loan supply side indicators. Unfortunately, we cannot examine this phenomenon more deeply because our dataset does not include any supply side information concerning loans. See Crook (2003) for an analysis also including supply side factors.

2.4 International Comparison of selected OECD countries

For international comparison we use data from the Luxembourg Wealth Study⁵ (LWS), which consist of ex-post harmonized international microeconomic datasets on wealth. Figure 2.2 shows the participation rate for the debt market in selected countries. In Austria, roughly 41% of all households hold either real estate debt or consumer debt. Compared to other countries, Austria's rate is rather low. In Sweden and the US in particular, the participation rate is relatively high. Germany is closest to Austria, but one ought to bare in mind that the

⁵<http://www.lisdatacenter.org/>

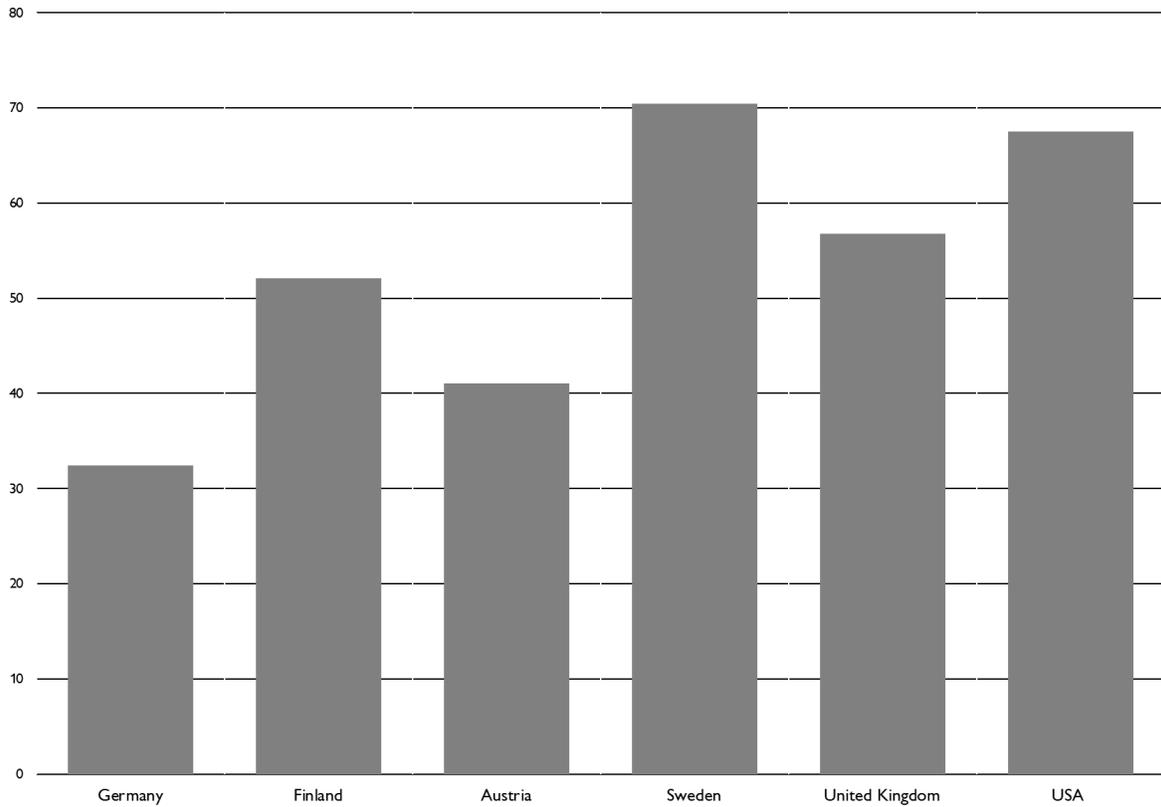


Figure 2.2: LWS, <http://www.lisdatacenter.org/>, Percentage ratio of households holding either real estate or consumer debt

rate for Germany is biased downwards, because the German dataset only reports debt values above 2500 Euros.

2.4.1 Empirical Results for Selected OECD Countries

In this section, we look at correlations between debt, income and financial wealth using LWS microdata for Germany, Finland, Austria, Sweden, the United Kingdom and the United States.

Total debt⁶ of a household $h = 1, \dots, H$, d_h consists of (secured) mortgage debt and consumer debt. Consumer debt itself consists of debt for the financing of vehicles, education, informal debt, and debt to consume other goods. We measure Total Financial Assets⁷ a_h

⁶LWS-Variable TD (Total Debt).

⁷LWS-Variable TFA (Total Financial Assets), which consists of *deposit accounts, total bonds, stocks*

and disposable household income⁸ w_h . All values are non-negative. As defined in (2.4.1), \tilde{d}_h includes only those households which hold positive amounts of debt. Note that with regard to total financial assets and household income we do not exclude values for which the logarithmic function is not defined as in (2.4.1), but include as many observations as possible using (2.4.2) instead. Analogous to (2.4.2), we treat \tilde{d}'_h in table 2.6⁹.

$$\tilde{d}_h \equiv \ln(d_h) \quad \forall d_h > 0, \quad (2.4.1)$$

total debt \tilde{d}_h , total financial assets \tilde{a}_h and disposable income \tilde{w}_h

$$\tilde{a}_h \equiv \begin{cases} \ln(a_h) & \text{for } a_h > 0, \\ 0 & \text{for } a_h = 0, \end{cases} \quad \tilde{w}_h \equiv \begin{cases} \ln(w_h) & \text{for } w_h > 0, \\ 0 & \text{for } w_h = 0. \end{cases} \quad (2.4.2)$$

Table (2.6) shows the correlation (Pearson) between total debt \tilde{d}_h , total financial assets \tilde{a}_h and disposable income \tilde{w}_h of households in the selected countries. In all cases, the correlation between debt and income is stronger than that between debt and financial assets. Both are significant at the 1% significance level. The second column \tilde{d}'_h includes households holding no debt. For all countries, disposable income stays positively and significantly related to debt. For Germany, the United Kingdom and the USA, this is also the case for financial assets. In Finland, the relationship between financial assets and debt turns to be significantly negative and for Austria and Sweden it is statistically insignificant. This is a first hint for income and financial assets being important determinants in the selection mechanism of holding debt or not. Indeed, Beer and Schürz (2007) conducted a logit regression and found that higher financial assets leads to a lower probability of holding consumer debt, but their results for housing debt were unclear.

and *mutual funds*

⁸LWS-Variable *lis_dpi* (Luxembourg Income Study Disposable Income), but exclude some negative values of disposable income for the USA, the United Kingdom and Sweden

⁹No dataset contains a $d_h, w_h, a_h \in (0, 1)$.

Table 2.6: Correlations of total debt with disposable income and total financial assets

	\tilde{d}_h	\tilde{d}'_h (analogous to definition (2.4.2))
Germany		
\tilde{a}_h	0.20***	0.03***
\tilde{w}_h	0.44***	0.39***
Finland		
\tilde{a}_h	0.13***	-0.07***
\tilde{w}_h	0.36***	0.29***
Austria		
\tilde{a}_h	0.30***	-0.02
Sweden		
\tilde{a}_h	0.25***	0.00
\tilde{w}_h	0.40***	0.45***
United Kingdom		
\tilde{a}_h	0.28***	0.09***
\tilde{w}_h	0.35***	0.37***
USA		
\tilde{a}_h	0.29***	0.26***
\tilde{w}_h	0.47***	0.42***
* * * =, ** =, * = denotes significance on a 1%, 5%, 10% level		

With respect to the perspective on all households, the positive correlation between income and debt as well as wealth and debt could be explained by borrowing constraints for less wealthy households but this share of households should not be large enough to explain this phenomenon in the data. Another possibility would be that variables other than financial assets or income drive the relation. Even in this case, the positive correlation documents a clearly suboptimal allocation in household portfolios. At the same time, the fact that higher amounts of debt are held by households with higher amounts of financial assets implies lower risks for financial stability than one could expect after perceiving decades of rising household debt on an aggregate level.

To control for other possible determinants of the selection into debt process and the amount of debt being held, we will turn to econometrics. The main goal is to check if the positive correlation will remain if we control for other household characteristics and the selection process of holding debt or not.

2.5 Empirical Results for Austria

Similar analyses were done by Crook and Hochguertler (2002) as well as Crook (2003) for OECD countries, Magri (2002) for Italy or Kohler et al. (2004) for Australia. To analyse the factors of debt holding in Austria, we use data from the Survey on Financial Household Wealth 2004, which is also included in the Luxembourg Wealth Study data. The correlation of total debt and total financial assets is positive and significant (see table 2.6). We see a strong relation between household total debt and financial assets. To compare the distributions of total financial assets and total debt, we produce kernel density estimates of \tilde{d}_h and \tilde{a}_h (see figure 2.3) for all indebted households¹⁰. To be able to compare, we present figure 2.4, which shows a similar kernel density estimation from Brown and Taylor (2005) for the USA¹¹. Figure 2.3 clearly shows that nearly all indebted households in Austria hold non-negligible amounts of financial assets, while in the USA there exists a non-negligible part of indebted households without any (or very small amounts of) financial assets. The distribution of total debt in the USA clearly has stronger weight to the right of the distribution than the distribution of financial assets. The absolute values of debt and wealth cannot be compared across the countries, because the data are collected at different points in time, the countries operate in different currencies, and they have different pension systems. However, we can say that the figures clearly show that indebted households in the USA optimize their portfolios more effectively, but that makes them more vulnerable to macroeconomic shocks.

In the case of a shock on the financial capability of serving debt payments, as a negative shock on income, unemployment, higher interest rates, or exchange rate or repayment vehicle shocks in the case of foreign currency loans, households need other resources to

¹⁰In this section we use total financial assets including life insurance, but results are robust with and without their inclusion

¹¹The estimates of Brown and Taylor 2005 use only indebted households and are based on data of the *Panel Study of Income Dynamics*, which is also part of the Luxembourg Wealth Study dataset

serve payments (see Albacete and Fessler 2010). If enough financial wealth is available they can draw from there financial wealth before they loose their real estate wealth, which serves as collateral in most cases. As fire-sales of real estate property due to macroeconomic shocks often lead to declining prices in real estate markets the value of the collateral might in the end only pay back a small share of outstanding debt.

To further analyze the relationship between total financial assets and total debt, our next step is to estimate a simple OLS Regression,

$$\tilde{d}_h = \alpha + \beta \tilde{a}_h + \gamma' W_h + \delta' X_h + u_h. \quad (2.5.1)$$

Our dependent variable is total debt \tilde{d}_h , made up of housing and consumer debt. This leaves us 991 observations holding positive amounts of debt in the Austrian dataset (1565 households without any debt are excluded from the analysis). α is a constant. The critical independent variables to analyze the relation between total debt and financial assets and disposable income are the logarithm of total financial assets \tilde{a}_h and a vector of three dummy variables W_h for increasing disposable income classes. The lowest disposable income class (up to 749 Euro) is used as reference class. X_h is a vector of socioeconomic control variables of the household or the household head h . u_h denotes a normally distributed error term with mean zero and variance σ^2 . The socioeconomic control vector includes age, age squared and dummy variables for the education of the household head¹² and a dummy for entrepreneurial-households¹³. We also include a dummy variable for households which have inherited money, and a variable indicating the number of members in the household.

We expect positive signs and increasing coefficients for the increasing disposable

¹²The lowest educational class is excluded as a reference category

¹³Entrepreneurial households are those in which at least one person holds not publicly traded business equity.

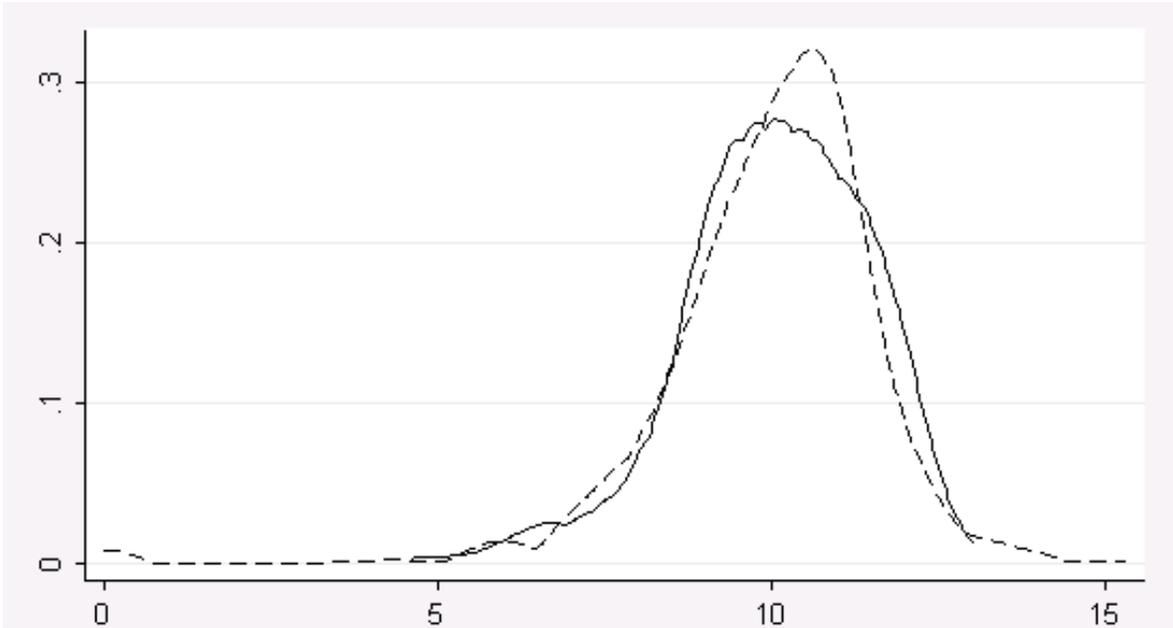


Figure 2.3: Kernel Denisty Estimation - Austria : — Total debt - - - Total financial assets

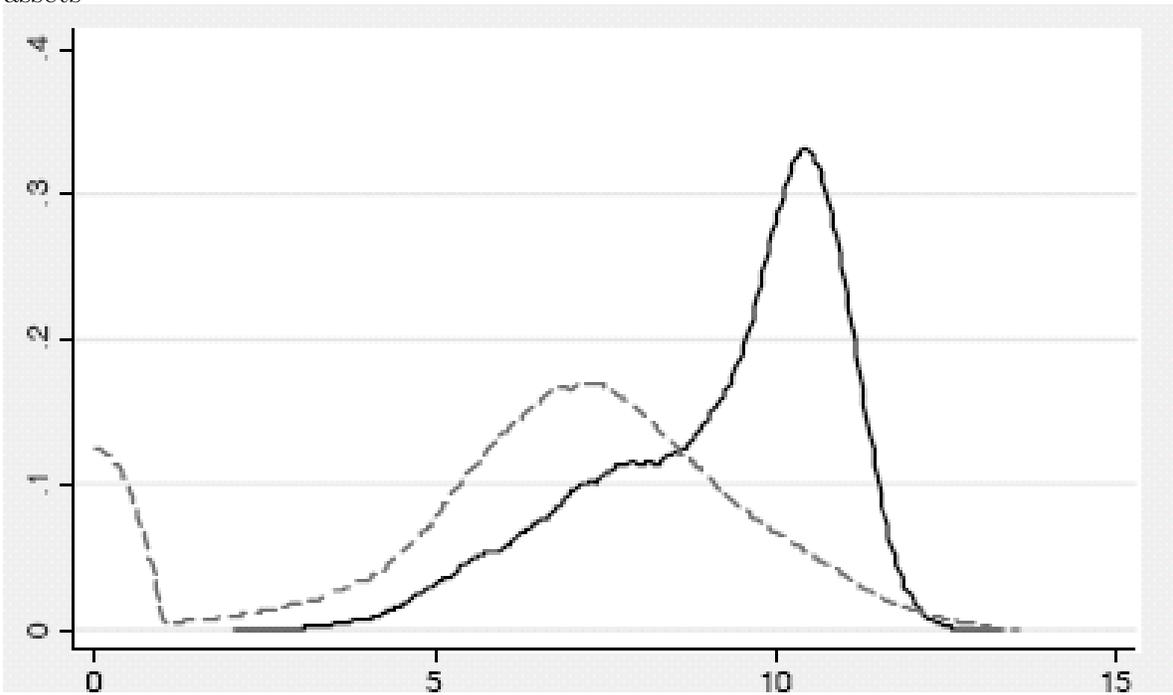


Figure 2.4: Kernel Denisty Estimation - USA: — Total debt - - - Total financial assets (Brown and Talyor 2005, Figure 9)

income dummies, as disposable income should be a good indicator for the ability to pay back debt. We also expect the same relation for the increasing educational dummies, because education should be a good indicator for the future development of an income stream of a household. We further expect a positive sign of the coefficient for the number of people living in a household, because this should be an indicator for the need of living space of a household and housing debt dominates consumer debt. For age of the household head we expect a negative sign, because the estimation just includes households with positive debt and we know that people in their thirties start to hold high amounts of debt financing their home ownership and older households should have already paid back most of their housing debt. The sign of the entrepreneur dummy and the inheritance dummy is not unambiguous. One argument could be that an entrepreneur is able to accumulate more debt because s/he can use the interest in business equity as security. On the other hand, an inheritance shifts the budget constraint to the right and should therefore lead to less demand of debt.

Because of the fact that we include only households with strictly positive amounts of total debt and the participation in debt market is very likely to be non-random (s. Beer and Schürz 2007), our OLS estimation could be exposed to *sample selection bias*. In order to control for that we estimate equation (2.5.1) as a Heckman Selection Model of the following form:

$$t_h = Z_h \delta + \mu_h \tag{2.5.2}$$

$$\hat{d}_h = X_h \beta + \varepsilon_h \tag{2.5.3}$$

$$t_h^* = \begin{cases} 1 & \text{if } t_h > 0 \\ 0 & \text{otherwise} \end{cases} \quad (2.5.4)$$

$$\hat{d}_h^* = \begin{cases} \hat{d}_h & \text{if } t_h^* = 1 \\ n.a. & \text{otherwise} \end{cases} \quad (2.5.5)$$

We estimate a Heckman selection model via the maximum likelihood method, where both equations are estimated simultaneously¹⁴. The model uses two latent variables, which are not observed, \hat{d}_h and t_h , which are dependent on two vectors of observable independent variables. Those are X_h and Z_h , where X_h consists of the same variables which we use in the OLS regression. For technical reasons the model requires Z_h to include additional variables, which influence the selection process (the 1st stage of the model)¹⁵. In our case the dummy variable for living in Vienna or not and the married/partner variable indicating that a spouse or partner lives in the household are candidates for additional first-stage variables. The variable \hat{d}_h defines the latent amount of debt and t_h the latent debt participation, i.e. holding debt or not. We observe t_h^* as an indicator for $t_h > 0$ and $\hat{d}_h^* = \hat{d}_h$, but only if $t_h^* = 1$. The error terms μ_h and ε_h are assumed to be jointly normally distributed with covariance $\rho\sigma_\varepsilon$. Put simply, (2.5.2) is a selection equation, explaining the decision to hold debt or not. Equation (2.5.3) is an equation explaining the amount of debt which is hold. If and only if a household holds debt, i.e. $t_h^* = 1$, the latent amount of debt is realised, i.e. $\hat{d}_h^* = \hat{d}_h$.

The difference between the Heckman selection model and a simple OLS estimation

¹⁴This type of model is also known as *Tobit Type Two Model*. For details on that type of model see e.g. Vella 1998.

¹⁵This requirement is under debate concerning the maximum likelihood estimation procedure we use here (see Vella 1998)

is that the OLS estimation includes only households with positive debt, whereas the Heckman selection model uses all households to estimate the desired latent amount of debt. The selection process is modelled explicitly in the first stage.

Our main concern lies on the households which hold debt and not on the desired levels of the others, i.e. the expected value $E(\hat{d}_h | Z_h \delta > 0)$ and the corresponding marginal effects of the independent variables on the amount of debt they hold. The advantage of the Heckman approach is that we can control for the selection process which could lead to biased OLS estimates, because households who decide to hold debt could be systematically different from households who do not hold debt (not only in terms of possible credit constraints). Marginal effects can be calculated from the *Heckman selection model* and are presented in table 2.7 jointly with the OLS results and the first stage of the Heckman selection model. The results show, that the positive and - at least at a 10% level - significant effect of total financial assets on the amount of debt is robust to controlling for the selection process. Note that marginal effects corresponding to dummy variables are effects from a discrete increase from 0 to 1.

Against our expectations, the hypothesis of a negative relation between total financial assets and the amount of debt is also rejected if one controls for a large number of socioeconomic characteristics. This result is robust for estimators which control for the selection process and, in this group, for resulting estimates corresponding to the total population (i.e. including not realised desired amounts of debt) as well as corresponding just to the households which hold debt (which are shown in table 2.7).

Besides the fact that a negative relationship between total financial assets and debt can be rejected, the signs of the coefficients for the other variables have a meaningful

interpretation. All significant coefficients of the OLS estimation have the same positive sign as the corresponding coefficients in the second stage of the Heckman selection model. In the OLS Model as well as for the marginal effects in the Heckman selection model (conditional on holding debt), all increasing income classes enter with increasing positive coefficients. The same is true for increasing educational classes, with the only exception of the highest educational class which has a marginally smaller coefficient than the second highest class, which could be explained by a later entrance in the labour market for people with tertiary education. The difference in the effects between these two groups is not statistically significant. The number of persons living in the household is as expected, with a positive sign. The entrepreneur- and the inheritance-dummy variables also have positive signs: households which have already inherited something seem to hold - *ceteris paribus* - more debt. The effects of age and age squared are not significantly different from zero in the second stage.

The *Heckman Selection Model* allows us to interpret the marginal effects of certain variables on the amount of debt, which just enter the selection equation. The Vienna dummy variable has a significant negative effect. One could argue that in Vienna, there are many more households renting apartments and who therefore do not need real estate debt.

The selection equation of the *Heckman selection model* shows that total financial wealth has a significant negative impact on holding debt. The probability that a household holds debt decreases with increasing total financial assets. Income has a significant positive impact on holding debt, which is the same as in the second stage of the model and a logit model estimated by Beer and Schürz (2007) on the same dataset. The same is true for education, but the education level is only significant (at the 10% level) for the lowest (included) educational class. Age and age squared enter significantly and

with the expected signs, age positive and age squared negative. That reflects the fact that debt participation reaches its maximum in the age class 30-39 and decreases for older household heads. Also consistent with this fact is that age and age squared enters the second stage (for marginal effects including all households which are not shown the coefficient is significant) in the opposite direction, because debt should decrease with age because people are repaying their debt.

2.6 Conclusions

Participation in real estate debt increases with higher income and higher financial assets. Participation in consumer debt increases with higher income but decreases with total financial assets. Housing debt dominates consumer debt with respect to participation and amounts of debt hold in this debt class.

The probability of holding debt (real estate and/or consumer debt) decreases - as one would expect - with higher total financial assets, once one controls for other possible determinants as income, education, age, and the number of people living in a household. For the group of households holding positive amounts of debt this negative relationship of total financial assets and total debt is - surprisingly - rejected. The more total financial assets an indebted household holds the higher the amount of debt - *ceteris paribus* - it holds. This is a strong hint for sub-optimal allocations in household portfolios of indebted households which goes a long with risks the households and the society as a whole unnecessarily bears. A mechanism of substitution between total financial assets and debt is rejected. Besides the fact that aggregate data is not useful to illustrate and analyse the resilience of households towards certain shocks, we find that the amount of debt a certain household holds is not enough to say anything about

Table 2.7: Ordinary Least Squares Regression, Heckman Selection Model (Full Maximum Likelihood), Austria, Total debt

	OLS	Heckman Selection Model	
		1. Stage (Selection)	2. Stage (Amount Estimation)
$E(\hat{d}_h Z_h \delta > 0)$			9.919
	Coefficient	Coefficient	for $t_h > 0$ Marginal Effects $\partial \hat{d}_h / \partial x$
Total Financial Wealth	0.095** (0.458)	-0.160*** (0.023)	0.074* (0.042)
monthly Net Income			
Euro 750.- bis Euro 1,349.-	0.858* (0.516)	0.540** (0.222)	0.619 (0.535)
Euro 1,350.- bis Euro 2,249.-	1.169** (0.511)	0.853*** (0.222)	0.906* (0.543)
Euro 2,250.- bis Euro 2,999.-	1.305** (0.519)	1.058*** (0.232)	1.021* (0.543)
Euro 3,000.- und mehr	1.476*** (0.524)	1.373*** (0.239)	1.167** (0.547)
Education			
Apprenticeship, Vocational School	0.446** (0.181)	0.190* (0.109)	0.396** (0.184)
Grammar School, Higher Vocational School	0.759*** (0.191)	0.114 (0.119)	0.687*** (0.190)
Fachhochschule, University	0.622** (0.220)	0.145 (0.132)	0.619** (0.206)
Number of Persons	0.101** (0.355)	0.080** (0.031)	0.085** (0.038)
Entrepreneur	0.288 (0.191)	0.152 (0.139)	0.302* (0.172)
Inheritance	0.250** (0.097)	-0.004 (0.067)	0.223** (0.093)
Age (Household Head)	0.000 (0.024)	0.057** (0.018)	-0.019 (0.023)
squared Age (Household Head)	0.000 (0.000)	-0.001*** (0.000)	0.000 (0.000)
Constant	7.101*** (0.796)	-0.486 (0.455)	
Variables from Selection Stage			
Vienna		-0.279*** (0.053)	-0.283*** (0.061)
married/partner		-0.086 (0.077)	-0.086 (0.077)
R-Squared	0.19		
LR χ^2 [p-value]		-1763.572 [0.000]	
Wald, $H_0 : \rho = 0$ [p-value]		74.29 [0.000]	
censored n		1565	
n	991	991	

the vulnerability of the household. Indebtedness of households needs to be analysed jointly with real- and financial assets of households. Given the empirical evidence of our analysis, it is quite likely that households with lower amounts of debt (consumer debt) which are neither covered by real- or financial assets could be more vulnerable than households with higher debt levels. On the other hand, that implies that household debt imposes a relatively small risk for financial stability, because the relatively low amounts of (consumption) debt which could trouble the households do not aggregate to amounts which could trouble banks.

2.7 Literature

Aizcorbe, Ana M., Arthur B. Kennickell and Kevin B. Moore. 2003. Recent Changes in U.S. Family Finances: Evidence from the 1998 and 2001 Survey of Consumer Finances, Federal Reserve Bulletin, January 2003.

Albacete, Nicolas and Pirmin Fessler. 2010. Stress Testing Austrian Households, Financial Market Stability Report 19, p.72-91. Oesterreichische Nationalbank (Austrian Central Bank).

Barnes, Sebastian and Garry Young. 2003. The rise in US household debt: assessing its causes and sustainability, Bank of England Working Paper, No. 206

Barwell, Richard, Orla May and Silvia Pezzini. 2006. The distribution of assets, income and liabilities across UK households: results from the 2005 NMG research survey, Bank of England Quarterly Bulletin, Spring

Beer, Christian, Peter Mooslechner, Martin Schürz und Karin Wagner. 2006. Geldvermögen privater Haushalte in Österreich: Eine Analyse auf Basis von Mikrodaten. In: Geldpolitik & Wirtschaft Q2/06. OeNB.

Beer, Christian and Martin Schürz. 2007. Charakteristika der Verschuldung der privaten Haushalte in Österreich: Ist die Verschuldung ein Risiko für die Finanzmarktstabilität?, Österreichische Nationalbank, Geldpolitik und Wirtschaft, No. 02/07

Brown, Sarah and Karl Taylor. 2005. Household Debt and Financial Assets: Evidence from Great Britain, Germany and the United States, University of Leicester, UK, Working

Cox, Pru, John Whitley and Peter Brierley. 2002. Financial pressures in the UK household sector: evidence from the British Household Panel Survey, Bank of England Quarterly Bulletin, Winter 2002.

Crook, Jonathan and Stefan Hochguertler. 2002. Household Debt and Credit Constraints: Evidence from OECD Countries, Credit Research Centre, University of Edinburgh, Working Paper Series No.05/02

Crook, Jonathan. 2003. The Demand and Supply for Household Debt: A Cross Country Comparison, Credit Research Centre, University of Edinburgh

Fessler, Pirmin and Peter Mooslechner. 2008. Arme Schuldner - Reiche Schuldner? Haushaltsverschuldung und Geldvermögen privater Haushalte auf Basis von Mikrodaten. In: Intervention, European Journal of Economics and Economic Policies, 5(1), 31-45.

Girouard, Nathalie, Mike Kennedy and Christophe André. 2006. Has the rise in debt made households more vulnerable?, OECD Economics Department Working Papers, No. 535

Herrala, R. 2006. Household indebtedness, Bank of Finland Bulletin, 2006-1

Kohler, Marion, Ellis Connelly and Kylie Smith. 2004. The composition and distribution of household assets and liabilities: evidence from the 2002 HILDA survey, Reserve Bank of Australia Bulletin, April

Laibson, David, Andrea Repetto and Jeremy Tobacman. 2000. A Debt Puzzle, Na-

tional Bureau of Economics Research, September, Working Paper 7879

Laibson, David. 1997. Golden eggs and hyperbolic discounting, Quarterly Journal of Economics, 112(2), pp. 443-477

Magri, Silvia. 2002. Italian Households' Debt: Determinants of Demand and Supply, Banca d'Italia, No. 454

Vella, Francis. 1998. Estimating Models with Sample Selection Bias: A Survey, The Journal of Human Resources, Vol. 33, No. 1.

Chapter 3

The Role of Data Production in Survey Analysis

3.1 Introduction

A large amount of economic research is based on analysis of observed data; this paper discusses some aspects of the relationship and disconnect between data collection and data analysis. It is usually the case that the data collectors are not the same individuals as those analysing the data, and in nearly all cases, collected data is heavily influenced by the individuals gathering the data. This fact is quite obvious for survey data, because survey data often comes from interviewers who directly gather information from respondents.

Nevertheless the same is true for so called register data. Also for register data it is true that either individuals gather information about a firm or institution, or firms or institutions report information they themselves gathered. Other individuals are compiling the data before they are used. All the way their are many possible sources able to influence data analysis in the end.

Therefore in any case it is important to know how data are collected - and how the process

of collection influences the data - in order to do reliable data analysis. The main objective of this paper is to provide guidance on working with survey data.

3.2 How survey production can drive the results of survey analysis

Survey data as a source for descriptive statistics and increasingly microeconomic analyses is heavily used in economics. However, in most cases, the collection and compilation of survey data is mostly done by statisticians while the analyses (especially microeconomics) is conducted (in the field of economics) by economists (microeconomists) who are not involved (in most cases) and not even familiar (in many cases) with the process of gathering the data.

This division between data compilers and data analysts is important, because the process of gathering the data can heavily influence statistical inference. In order to provide valid results, a microeconomist needs to be aware of and take into account the process of data collection. In this section of the paper, we provide a non-technical overview of several steps in the data collection and generation processes which can influence statistical inference. For the sake of simplicity, we concentrate on household survey data, but most of the arguments hold for all various types of survey data (e.g. data on firm level, personal consumer surveys, etc.)

The main problem of statistical analysis with regard to some population at hand is a simple one. We want to be able to make true statements about the population, but in order to do so, we would need information on all the elements of the population. Most of the times this is not feasible because it would be too demanding - in terms of time and costs - to gather all the necessary information we need from every member of the population. The solution to this problem - at least to a certain degree - is to survey the population in a representative

way. That means that we collect the necessary information for a subset of the population, the so-called sample, which consists only of a very small fraction of the population. If we do that according to some rules, it can be shown that we will be able to make true statements about the whole population which are only based on information on the subset of the population (sample). Because we do not have all the information on all of the elements of the population, this statement comes with some degree of uncertainty, which is a known degree of uncertainty, so long as everything is done correctly.

In order to make correct statements about the population on the basis of the sample and to calculate the correct degree of uncertainty of these statements, two necessary conditions have to be fulfilled. The first is proper collection of the data, which should include good documentation of all its steps. The second is the use of a number of adjustments and use of the correct statistical methods to account for a given process of data gathering. In this section, we review critical components of these two steps - the sample frame, sampling design, survey mode, interviewing, paradata, weights, and imputation - to discuss how survey data ought to be collected and compiled.

3.2.1 Sampling frame and sampling design

The sampling frame is a set of elements which optimally represents all of the elements in the population which a researcher wants to make statements about. The sampling design includes all the rules regarding how these elements are drawn into the sample.

Both the sampling frame and the sampling design are important and can have serious impacts on the results a researcher finds when analysing the sample. In general one has to distinguish between two different cases (most of the time only the second is relevant):

1. The researcher is only interested in making statements about the sample at hand, i.e.

statements about the subset of elements actually surveyed (e.g. household taking part in a certain survey).

2. The researcher is interested in making statements about the whole population by analysing only the sample, i.e. the subset of elements actually surveyed.

Most of the time, researchers are interested in making statements about the whole population and not only the sample. In this case, the sampling frame and sampling design is relevant in order to be able to make true statements with a known uncertainty. This will typically be a statement like:

“The average household income of the household population is y and if we would repeat this survey using the same sampling frame and design indefinitely, 95% of the resulting y s would lie in between x and z .”

The more elements the sample includes, i.e. the more elements are drawn from the target population into the sample, the narrower this 95% confidence interval will be. In other words, a larger sample leads to a more precise estimate of y . How to calculate this uncertainty is another question and there are different possibilities, which have all their pros and cons. But what is more important for us is that the estimate as well as its precision is not only influenced by the different methods of calculation but possibly (and typically to a larger extent) also by the sampling frame and sampling design.

Sample frame In the case of household surveys, which refer to the population of households in a certain geographical area (usually a country), many different types of sampling frames are used. Some are based on lists of households (or individuals), some are based on certain areas and some are based on other entities like buildings, telephone numbers or mailing codes.

An optimal frame includes units which represent all elements of the population without double counts, i.e. all elements are included exactly once in the sampling frame. In most cases, an optimal sampling frame is not feasible but there will be some over and/or under coverage of certain parts of the population. In order to minimize possible bias of the sample, the first goal is to minimize under- and/or overcoverage.

A typical example of a problem with coverage would be a household survey which uses a list of telephone numbers as sampling frame. In this case, larger households (those with more household members) which typically have more telephone numbers would be overrepresented, because the probability of one of them being drawn into the sample is higher than for smaller households. On the other hand, households without any telephone numbers (or whose telephone numbers are not included in the list) are excluded. These households do not enter the sample because they are not covered by the sampling frame. An attempt to minimize coverage problems of this type would be to try to identify numbers which refer to the same dwelling (e.g. using addresses linked to the telephone numbers if available) and cluster them in the process of drawing households out of the list of telephone numbers to reduce the problem of overcoverage.

Even in relatively good sampling frames, sampling households from official dwelling lists (e.g. Austrian Melderegister or lists of private dwellings from a post office) still lead to over- as well as undercoverage problems. For that reason typically “institutionalized households” - such as people living in prisons, homes for the elderly, or members of church and military institutions - are excluded. The homeless and the super-rich are also hard to reach in typical household surveys, giving rise to the so-called middle-class bias of such surveys. The US Survey of Consumer Finances (SCF) e.g. excludes - due to confidentiality reasons - the Forbes-Top-400 high wealth individuals from the sample. The EU-SILC in Austria, for example, restricts its sample frame to households who live at a dwelling with official “Hauptwohnsitz”-Status (registered main-residence) in the Austrian “Melderegister”, which e.g. excludes many students living and working in Vienna, but who have their official “Hauptwohnsitz” still at their

parents' main residence. It also excludes any people who do not want to register a "Hauptwohnsitz" but prefer to live in their "Nebenwohnsitz" for most of the year. This illustrates how important it is how a household is defined, and to control which target population is covered by the sample. Therefore, the definition of the elements which are a part of the sampling frame is very important to understand what type of population the researcher can or cannot make statements about in the end.

Of course sample-frame issues become especially important in cases as access-panels samples or experiments, when the sample is selected in a non-probabilistic way. In that case it is very difficult to give estimates and corresponding standard errors a meaningful interpretation with relation to a certain target population. This leads us to the sampling design which defines how the sample is drawn from a given sample frame.

Sampling design The sampling design includes all the rules regarding how elements in the sample frame are selected into the sample. Given an optimal sample frame (all elements of the population are included exactly once), the most efficient way to draw elements into the sample would be a simple random sampling scheme where all elements are selected with equal probability of 1 divided by the number of elements. Again, the more elements drawn into the sample, the more precise the resulting estimates (statements) would be. Note that given some ex-ante knowledge of the target population there might be good reasons to oversample certain parts of the target population. As examples it makes sense to oversample (i) parts of the population which are known to have higher non-response rates than others or (ii) parts of the population which hold a disproportional part of a certain item one wants to measure, i.e. parts of the population which contribute more to the overall variance of a certain item than other parts.

In practice, there are many other constraints that hinder the sample design from being completely simple random sampling, and instead consist of a sampling process using a number

of stages. In a staged sampling process, elements are grouped into sets and those sets can be grouped again (this can go on many times) and then elements are taken from the subgroups. A typical example would be a household survey, in which in the first stage municipalities are drawn, then in the second stage streets or blocks are drawn, and in a third stage households are drawn. Of course such a type of sampling reduces costs but (in general) it increases variance, because of the clustered nature of the resulting sample and the fact that elements grouped (e.g. regionally) in a certain set are probably more homogenous than elements in general. For an overview of different sampling possibilities see Kish (1995).

Why should the analyst care about the sample frame and sampling design?

First, the sample frame implies to which target population statements resulting from survey analysis refer to. For example, typically sample frames for household surveys exclude “institutionalized households”. Therefore statements resulting from survey analysis using those data are not including those households. Or in the case of EU-SILC in Austria, which only samples households which are living in a dwelling where a “Hauptwohnsitz” is registered, all resulting estimates are not referring to all other households living in dwellings where no “Hauptwohnsitz” is registered.

Again, data gathered in experimental studies are a rather extreme case very well illustrating the problem. An experiment conducted with voluntary participants at a certain university is not representing the overall population and very likely even not the student population or the student population at the university where the experiment took place. If a well defined sample frame as well as a well defined sampling design those not exist it is impossible to relate resulting estimates to a certain target population. Instead statements can only be drawn with regard to the sample at hand and standard errors of estimates have no clear interpretation with regard to a certain target population as long as the selection process is not controlled via the sample frame and sampling design.

Therefore it is of utmost importance to know details about the sampling frame and design used to be able to relate the resulting estimates to a well defined target population and to give the corresponding standard errors a meaningful interpretation.

Secondly, the sampling design is important in terms of weighting. Even if there is proper knowledge and documentation of the sample frame and sampling design this information needs to be taken into account correctly when analysing the data. Without controlling for the sample design, most estimates will be biased (see 3.2.5 for details). Design weights, which should give the inverse probability of a certain element being drawn in the sample, are usually calculated but seldom provided in research datasets. With design weights, one can correct for possible unequal selection probabilities. Unequal selection probabilities may be introduced for several reasons, the most important of which are cost reduction and/or oversampling of a certain group in order to obtain higher precision of that particular group. What is mostly included in research datasets are so called final weights, which should - at least in theory - also include design weights.

There is hardly any dispute in the literature concerning the use of survey weights when estimating population means or elasticities, and recent literature argues clearly in favour of using weights in regression analysis as well (i.e. doing WLS (Weighted Least Squares) instead of OLS (Ordinary Least Squares)), as soon as the sample is non-random and weights are given as the inverse sampling probability (Angrist and Pischke 2009; Faiella 2010).

3.2.2 Survey mode

The survey mode is the method of gathering data in a survey environment. There are different methods available and different methods are suited for different situations and goals. The main question is what quality of the gathered data is necessary in order to be able to conduct certain analyses. When the population of interest consists of private households (but also persons,

firms, etc.) there are several possibilities at hand. Depending on the complexity and amount of the data which should be gathered, some survey modes are more appropriate than others. In general there is a strong trade off between quality and cost. These days, mostly the following survey modes are used:

- Computer Assisted Personal Interviewing (CAPI).
- Computer Assisted Telephone Interviewing (CATI).
- (Audio) Computer Assisted Self-Interviewing ((A)CASI).
- Paper And Pencil Interviewing (PAPI).

For a long time, the most common survey method was PAPI. In PAPI methods, a paper-based questionnaire is sent to the household, and a certain member of the household is supposed to answer the questions and send the survey back. Problems with this type of data gathering are obvious, as many complex surveys are difficult for respondents to understand, and the program of the questionnaire should differ for different households. For example in a survey on housing, an owner-occupier household should answer questions concerning the primary residence the household owns whereas different questions are relevant for households who are tenants in their primary residence.

(A)CASI interviews are mostly web-based and have certain advantages compared to PAPI surveying, as one can program a path dependent questionnaire and can provide some help to the respondents when necessary. As there is no personal interaction taking place between an interviewer and a respondent, both (A)CASI and PAPI are not exposed to interviewer bias, which we will discuss later. On the other hand, a main problem is that non-response is in general much higher in self-administered surveys and problems of understanding the questionnaire can be very severe especially in complex surveys.

CATI as well as CAPI involve an interviewer asking the respondent the questions and guiding him/her through the questionnaire. The main advantage is that the interviewer, if trained properly, can be of enormous help to explain unclear questions and decreasing unit-as well as item-non response.

Especially for complex and demanding surveys, CAPI is considered to be the best choice. The interviewer is in personal contact with the respondent and can explain questions to him/her. Furthermore, supplemental material, such as paper based cards, can be used which help the interviewee to answer questions, especially if there is the need to choose from a large number of categories. Moreover, direct plausibility checks can be programmed and immediately executed and worked on during the interviews, leading to more consistent interviews and preventing misunderstandings. For these reasons, most complex scientific surveys today use the CAPI method (see Groves et al. 2004).

In EU-SILC some countries use only CAPI (e.g. Italy), some only self-administered PAPI (e.g. Germany), some only interviewer-based PAPI (e.g. Hungary) and some only CATI (e.g. Iceland) - but many of them use a mixture of the aforementioned possibilities. The literature is pretty clear about the fact that those methods have all a certain impact on the data gathered with the method at hand (see Groves et al. 2004; De Leeuw 2008)

Hauser (2007) clearly states the problems arising for the German EU-SILC dataset which is collected using only self-administered PAPI. That this leads to different results than a fully-fledged CAPI-Approach is obvious. Hauser mentions several points questioning the quality of the dataset and therefore especially its comparability to datasets collected in a different way. We mention just a selection of the many arguments here:

- Self administered interviews lead to higher non-response rate
- Sending interviews by post leads to severe problems when trying to resolve inconsisten-

cies ex-post via telephone; those households who can not be reached because they have no landlines are significantly different from the ones with landlines

- Self-administered complex surveys lead to many misunderstandings and wrong answers due to the misunderstanding. Those misunderstandings are non-random but correlated with education, income and many other variables in the survey
- Sample material is sent out in German, but there is no way to control for non-response due to language problems, because it is not known why a household did not answer: there could be no household at a certain address, the household might have denied participation, the household could have problems understanding the material, etc.

While many surveys use mixed mode strategies in the contact phase, the use of mixed modes in the response phase is less frequent. Nevertheless mode changes during a response phase are increasingly used especially when dealing with sensitive questions. In those cases the mode change occurs for all the respondents for certain questions. Mode changes in the response phase in the sense that for a part of the respondents a certain mode is used while for others a different mode are less common.

While in the first type of mode change in the response phase still all respondents answer all questions given the same mode as all the other respondents in the second type different respondents answer the same questions given different modes. The Literature calls this situation a mixed-mode design¹ and it is e.g. used in the Austrian EU-SILC survey (see De Leeuw 2008). In this case the question arises if the mode has an impact on response behaviour given a certain question.

Since the 2008 wave, EU-SILC in Austria uses CATI as well as CAPI as interviewer modes. Households which can not be reached via CATI or opt for CAPI instead of CATI are turned

¹Those can be split up into concurrent and sequential mixed mode designs. The main difference is that in the sequential design only for those who are not reachable with the first mode another mode is tried in a follow up phase, whereas in the concurrent mixed mode designs people can (also) opt for a certain mode

into CAPI cases. The selection bias of this method is obvious: for households which are easily reachable via telephone and are willing to provide information via telephone the information is gathered via CATI, whereas the information on the rest of the households, if they can be reached personally, is collected via CAPI. As there is no identifier included in the user (researcher) dataset (at least thus far for EU-SILC 2008) an analyst - even if s/he would like to - can not control for the influence of the survey mode. And more worrisome: even if s/he could, nobody knows which part of the difference is due to the mentioned selection bias and which is due to the survey mode, because the survey mode is not distributed randomly to households in the sample.

To illustrate the impact of the use of different interview modes we look at the Austrian EU-SILC wave 2008. Table 3.1 shows the differences in the income distribution between CAPI and CATI observations in the EU-SILC wave 2008. Note that we control for selection bias via coarsened exact matching on variables from the 2007 EU-SILC wave where all households were still interviewed with the same interview mode (CAPI). Fessler, Kasy and Lindner (2011) use a non-parametric 1 to n coarsened matching procedure based on all significant variables in a selection equation to balance the observations on a large set of covariates including income in 2007². Whereas the gini coefficient is 0.30 (0.31 in the unbalanced sample) in the CATI sample, it is 0.32 (0.34 in the unbalanced sample) in the CAPI sample. Furthermore they show that the use of CAPI significantly reduces item non-response for income questions as well as leads to higher values of income especially at the top of the income distribution.

All those facts have an impact on the resulting joint distributions of the variables. This is true for all different survey modes. Some things are more problematic to talk about on the phone than in person, some things can be difficult in a self administered survey, which could be explained by an interviewer on the phone and would be no problem at all in a personal interview.

²2199 out of 3772 Households on common support (430 out of 1652 covariate combinations)

Table 3.1: Differences in Gini Coefficients over modes with bootstrapped standard errors

	Bootstrapped Differences	Bootstrapped Differences: Balanced Observations
CAPI-CATI	0.029** (0.013)	0.018* (0.009)

Notes: This table shows the estimates for the difference in Gini Coefficients for the unbalanced and balanced datasets.

Moreover in some cases so called “output oriented” datasets are using different methods of data gathering when compiling a dataset. EU-SILC is a output oriented project meaning that there are a large amount of output variables defined and every participating country gathers the data to compile their EU-SILC dataset. Then EUROSTAT merges the dataset and researchers use the dataset as if it were one single dataset. The data gathering techniques are not harmonized over the different countries participating in EU-SILC, and this ends up being quite problematic. Lohmann (2011) illustrated the resulting problems for EU-SILC data using the relationship between employment, earnings and poverty as an example. He differentiates between four basic approaches to data collection varying across EU-SILC countries:

- Survey: income and employment information from the same source (personal interview)
- Survey/proxy: income and employment information from the same source (personal proxy interview, i.e. a person from a certain household gives the information about another person of the same household)
- Register: income and employment information from two different sources (register data and personal interview)
- Register/proxy: income and employment information from two different sources (register data and proxy interview)

The data gathering process is different in all different settings in the sense that it combines (or not) register data with personal interviews and proxy interviews when gathering certain information about individuals living in private households. Lohmann (2011) finds that there is a significant effect of the different types of data gathering on poverty rates by activity status and on the poverty risk of non-working and working population.

Why should the analyst care about the survey mode? In a perfect world, the survey mode would have no effect on the respondents' answers and the impact on unit- and item non-response could be perfectly corrected with statistical methods, and an analyst would therefore have no need to care about the survey mode. But in general, the situation an analyst is confronted with in the real world is different. The European Union Statistics on Income and Living Conditions (EU-SILC) is one of the best examples to illustrate the severity of problems arising from the use of different data gathering techniques and different survey modes.

In sum, if an analyst is interested in the question of being able to make true statements with a known uncertainty about a population at hand, s/he must first be interested in how the data was gathered. It is important that any differences in how samples were drawn and differences in the survey mode be noted in the data and available to the researcher. In the case of country comparisons and different methods of data gathering, the minimum one can do is to use country fixed effects, but which neither helps to identify the impacts of the different types of data gathering techniques - including different survey- and interviewing modes - over countries nor eliminates the problem of biased joint distributions due to different data gathering techniques - including different survey- and interviewing modes - inside one country dataset.

If identifiers for different survey modes, proxy interviewers or other differences are available

they should be controlled for in all types of regression analysis. If one has reason to believe that the impact of a certain data gathering technique might be different for different groups of households additionally interaction terms should be used to allow for heterogenous effects over the covariates used. Ignoring differences in data gathering might likely lead to severe bias.

3.2.3 Interviewing

Interviewers and interviewing itself is by far the most underestimated ingredient when it comes to data collection. Interviewer-bias, the influence interviewers have on the respondents answers, is well covered in the relevant literature (e.g. Groves et al. 2004). The use of a large number of interviewers is a way to reduce interviewer bias in the overall dataset. Nevertheless the influence of certain interviewer characteristics on certain responses hardly receives any attention in empirical analysis. The main reason is pretty straight forward because there are basically no research datasets available which cover information on the interviewers. Nevertheless, such information could influence the datasets and therefore research through different channels. The most important are for sure unit- and item non-response which are covered in sections 3.2.5 and 3.2.6. But also the content - values and outcomes - of the information gathered in an interview can be influenced by certain interviewer characteristics. There could be more trust with some types of interviewers and respondents than between others, some interviewers might probe in a less or more manipulative way than others, and some interviewers even induce certain behaviour just by their appearance or by making statements that induce ideas about the values of the interviewer to the respondent. All of these possibilities can lead to bias.

Why should the analyst care about interviewing? If there is information about the interviewers available in the research dataset, the analyst can try to control for certain characteristics when analysing the data. This should help to reduce interviewer-bias. Of course,

one has to be aware of the possible high correlation of interviewers with regional variables. As such information is almost never available, the analyst should at least gather information on how the data producers tried to reduce this type of bias. This should typically include detailed interviewer training, documentation about the number of interviewers and guidelines for their work, including material which was provided for a successful interviewing situation (handbooks, glossaries, etc.). The analyst should also check if the data producers gathered information about interviewers to use that in the process of weighting for non-response and dealing with item-non response in the imputation process.

3.2.4 Paradata

Paradata are all data, which are not part of the interview itself and therefore not provided via answers of respondents. Typically that might be information about the dwelling of a household like the surrounding area, the type building, the availability of a car-port or facts about the interview itself, like if a respondent used certain documents when answering questions or how many people were around. The larger the amount of paradata available which does not depend on the participation in the survey the better for controlling for unit non-response issues.

As we mentioned in section 3.2.1 it is of utmost importance to know which target population is represented in the sample and which parts might be under- or overrepresented. We will go into deeper detail concerning that issues in section 3.2.5. For now it is important to mention, that unfortunately most survey datasets only consist of that part of the sample which in the end participated in a certain survey. All observations which are not eligible or did not participate in the survey are usually not included in research datasets. Furthermore in many surveys no data at all is gathered about sampled households which did finally not participate in the survey, either because they did not want or could not be reached. Therefore typically only few information is available to model the process of unit non-response and calculate correct non-response weights (see 3.2.5).

Why should the analyst care about Paradata If the analysts goal is to make statements about the sampled target population s/he needs to care about the process of non-response. This is crucial because non-response is very unlikely randomly distributed in the sample and therefore the observations available to the analyst might misrepresent the original target population represented in the sample. For that reason unit non-response weights should be calculated (see section 3.2.5).

If the researcher dataset includes the whole sample (including non-respondents) and also includes paradata information an analyst can model the process of unit non-response. In most cases though, this crucial information is not included in research datasets and very often paradata - if gathered at all - is not gathered for all households, including the non-respondents. As most survey dataset are available only including so called final-weights, the only thing a thoughtful analyst can do is to check the documentation if non-response weights (adjustments) are included in the final weights. Furthermore it is important to know the actual non-response rate and how much paradata was available to produce non-response weights (adjustments) to get an idea about the degree to which the original target population is represented by the resulting dataset. Quality - with regard to the problem at hand - rises with decreasing non-response rate and with the amount of paradata available for all sampled elements, which was used in the non-response adjustment.

Of course the worst case scenario is one where no information at all can be provided about non-response. This happens when a survey had no fixed apriori sample, which means that data is gathered as long as it takes until a certain number of participating households is reached. In that case there is no possibility to tackle the problem of unit non-response in a sensible way.

3.2.5 Weights

Weights are part of most survey datasets. Many analysts have the idea that they are only reweighting the observations to fit certain population-characteristics (as e.g. the distributions

of age, gender, region or household size) and as one controls for those variables in a regression, they think they can forget about weights. If a survey is carried out poorly, that might be true. But for most scientific surveys today, that idea is wrong and can likely lead to serious bias in estimates. In state of the art scientific surveys a set of weights is produced in order (i) to correct for certain misrepresentations of the sample and (ii) to allow the correct estimation of variances even without sampling variables (like PSUs, SSUs, STRATA, etc.), which are - due to confidentiality issues - mostly not available in research datasets.

Misrepresentation, in the sense that the sample characteristics are not to an argueably degree representing the target population characteristics, arises due to different phenomena, of which the most important are:

- Sampling Design
- Erroneous inclusions
- Erroneous exclusions
- Non-Response

The sampling design leads to misrepresentation as soon as the probability of being drawn into the sample is not the same for all elements in the sample frame. That is very often the case in complex survey designs and is sometimes a goal in order to increase the precision of estimates for certain groups and/or items (oversampling). Nevertheless as long as the probability to be drawn into the sample is positive and known for every element in the sample frame the inverse probability is the correct design weight and is therefore already given apriori to correct for the resulting misrepresentation due to the sampling design.

Examples for errorneous inclusions are e.g. multiple (unknown) inclusions of households in the sample frame and for errorneous exclusion e.g. if more than one household is located in one dwelling or households are sampled via telephone adresses and households without telephones exist. Non-Response leads to misrepresentation as soon as it is not completely random.

Basically every dataset suffers from one or (mostly) more of the problems mentioned. If an analysis of the data at hand is restricted to making statements only about the subset (in the case of erroneous inclusion not a subset but including also elements not included in the target population) of the population covered in the dataset, those problems can be ignored and an analysis without reweighting, i.e. assuming equal weights for every observation is in order. If the goal is to make statements about the target population correct weights have to be considered in most cases.

Typically the final weights included in most surveys (should) consist of the following weights correcting for different misrepresentations:

- Design weights (correcting for unequal sampling probabilities)
- Coverage weights (correcting for under- or overcoverage in the frame)
- Non-response weights (correcting for the non-random part of the non-response process)
- Post-stratification weights (fitting the sample closer to distributions in the population; external data needed)

In most surveys, a so called final weight is provided (which includes in the optimal case all the above mentioned weighting steps). Given that weight as well as sampling information (Primary Sampling Units (PSU), Secondary Sampling Units (SSU), STRATA, etc.) one can calculate unbiased (in the sense of design unbiased not model unbiased as described by Faiella 2010) estimates as well as the correct variances, taking into account the sampling process. Many statistical software packages provide special packages to estimate the correct variances using the sampling information and linearization (however some commonly estimated parameters do not allow for standard error estimation via linearization - e.g. the median). But for many surveys, sampling information is - due to confidentiality issues - not provided in the research datasets. In order to be able to calculate the correct variances of the estimates, so-called replicate weights are included in at least some state of the art datasets. These are

calculated by the data producers and are based on numerical methods (mostly bootstrap and jackknife procedures) using the sampling information. Properly used, they allow for a correct variance estimation even in the absence of sampling information as STRATA, PSUs or SSUs.

When constructing the weights, data producers face a bias versus variance trade off. Their decisions can heavily influence later results. A typical example is the trimming of weights: Let us assume that for weighting reasons households were grouped with regard to some socioeconomic characteristics as well as sampling information and paradata available. As a next step these groups would be ordered -maybe by socioeconomic characteristics as well as paradata - according to their non-response rate. Let us assume that the group with the highest non-response rate has a significantly higher non-response rate than all the other groups in the sample, implying very high weights for this subgroup relative to other parts of the sample. A common practice if the weights (and therefore overall variance) would get “too high”³ is to merge the subgroup with the next closest subgroup and weight them together. Let us assume that the group with the highest non-response had a non-response rate of 80% - implying non-response weights of 5 - and the group with the second highest non-response rate had a non-response rate of 60% - implying non-response weights of 2.5. Merged together - assuming the second grouped included more individuals - the non response rate of the resulting group is 66% - implying non-response weights of 3. This pooling of different households of course implies a certain misrepresentation because two different groups for which non-response was different are now treated as one - more heterogeneous - group of households for which equal non-response weighting is applied. This of course results in less variance but higher bias implying a higher probability of statistical significant results when doing survey analysis (see 3.2.6 for a similar problem). In practice one has the choice to introduce as many groups as reasonably possible - given the amount of information available for the whole sample - minimizing the bias resulting from non-response but maximizing the variance introduced by this step. Or, calculating one single non-response rate and use that one for weighting to control

³The rules applied in that context, i.e. what is considered as “too high”, seem to be rather arbitrary.

for non-response, which of course means nothing else than not weighting for non-response at all, i.e. ignoring the non-response process - maximizing the bias resulting from non-response but minimizing the variance introduced by this step.

Why should the analyst care about survey weights? Most of the time researchers want to make statements about a certain population of interest, called the target population. If so, either reweighting or strong assumptions concerning most methods applied are in order. There is hardly any dispute in the literature concerning the use of survey weights when estimating population means or elasticities and recent literature argues clearly in favour of using weights also in regression analysis, i.e. doing WLS (Weighted Least Squares) instead of OLS (Ordinary Least Squares). The advocates of WLS claim that WLS parameters are more robust because they are model unbiased if the model is true and design consistent if it is not. If the OLS model is misspecified and predictors correlated to the response variable are omitted, OLS estimates may be biased and inconsistent (Faiella 2010).

Nevertheless under certain (pretty strong) assumptions⁴, the WLS estimate will be consistent for β for any vector of weights, including the OLS case of a vector of equal weights. Furthermore, if the error ε_i is homoskedastic the case of equal weights (OLS) is the most efficient one. Note that for the necessary assumption $\mathbb{E}(\varepsilon_i|x_i) = 0$ to be reasonable, the determinants of the sampling frame and non-response process need to be included in the controls (Cameron 2010).

An easy check is to do both OLS and WLS. If the estimates are significantly different from each other, that would indicate that $\mathbb{E}(\varepsilon_i|x_i) \neq 0$, and WLS is in order. The risk of doing OLS when indeed WLS would be in order is getting biased significant estimates due to model misspecification. The risk of using WLS if OLS is in order - in the sense that all necessary conditions are fulfilled - is that a unbiased estimate could be not significant even though it

⁴Namely that the DGP (Data Generating Process) is in fact the specified model $y_i = x_i'\beta + \varepsilon_i$ and sufficient controls are assumed to be added so that the error $\mathbb{E}(\varepsilon_i|x_i) = 0$

would be significant in a more efficient OLS specification with weights being equal.

Basically the problem is analogous to the question of using a random effects model instead of a fixed effects model: the random effects model is a parametrical restriction of the fixed effects model. In the same sense, the OLS model is a parametrical restriction of the WLS model, namely with the restriction of weights being equal for all units. In that spirit some argue in favor of using a kind of Hausman test. Under the null hypothesis, the estimates of OLS and WLS are not different and both WLS and OLS are consistent. In that case OLS, using equal weights, is more efficient and WLS would pointlessly increase the parameters standard errors (as the fixed effects model when indeed the random effects model is in order). Under the alternative hypothesis only the WLS estimator is consistent whereas OLS estimates are inconsistent.

Theoretically, a broader range of possible tests in order to identify if design and non-response is ignorable or not, i.e. if OLS is consistent or one should work instead with WLS, is available. The problem in most cases is, that due to confidentiality and other constraints most researchers are not able to perform those tests. In many cases only final weights are delivered, but not design, non-response, post-stratification weights separately. Very often information on strata and other necessary sampling information is not included in the datasets available to researchers. To sum up, to be on the safe (always consistent) side, WLS is in order. If one is dealing with efficiency problems and therefore wants to use OLS, both resulting estimates should at least be compared. If they are significantly different, the OLS estimator is biased. All of this applies to all other models assuming equal weights.

Whereas for means etc. the effects are obvious table 3.2 shows the impact of the use of weights versus not using weights in an example of regression analysis. While the signs of the estimates are not changed by the use of survey weights (which is of course theoretically possible), coefficients and standard errors differ even if no imputation but listwise deletion is used (as in columns OLS1 and OLS2 of table 3.2. In the case of using simple imputation

and therefore controlling for item non-response the impact of using weights is already very pronounced (see columns OLS3 and OLS5 of table 3.2), i.e. 3 out of 7 coefficients turn out to be not statistically different from zero on a 5% significance level as soon as weights are used. For the estimation using multiple imputations given in column OLS6 (which can not be reasonably done without weights), which is the only one which provides the correct standard errors including the uncertainty introduced by imputations the effect is even stronger (see section 3.2.6 for more details).

Table 3.2: Household Survey on Housing Wealth 2008; OLS Regressions of household real estate wealth on household characteristics illustrating the use of using weights and/or single/multiple imputations,

	OLS1	OLS2	OLS3	OLS4	OLS5	OLS6
Vienna	-0.358*** (0.122)	-0.308*** (0.106)	-0.359 (0.251)	-0.275*** (0.087)	-0.334 (0.296)	-0.269 (0.318)
Inheritance	0.768*** (0.154)	0.672*** (0.173)	1.944*** (0.253)	0.543*** (0.133)	2.092*** (0.426)	1.964*** (0.440)
No. of adults	0.427*** (0.062)	0.410*** (0.084)	0.410*** (0.125)	0.327*** (0.066)	0.396*** (0.099)	0.467*** (0.112)
Age	0.085*** (0.017)	0.081*** (0.013)	0.099*** (0.035)	0.072*** (0.011)	0.065 (0.045)	0.067 (0.047)
Age squared	-0.001*** (0.000)	-0.001*** (0.000)	-0.001** (0.000)	-0.001*** (0.000)	-0.000 (0.000)	-0.000 (0.000)
Female	0.107 (0.097)	0.095 (0.096)	0.204 (0.200)	0.090 (0.079)	0.168 (0.194)	0.183 (0.196)
Medium education	0.500*** (0.140)	0.500*** (0.117)	0.660** (0.288)	0.413*** (0.096)	0.574** (0.228)	0.528** (0.254)
High education	0.619*** (0.160)	0.616*** (0.151)	0.684** (0.331)	0.520*** (0.126)	0.641** (0.281)	0.589* (0.333)
Constant	-2.692*** (0.411)	-2.561*** (0.326)	-2.780*** (0.865)	-2.093*** (0.269)	-1.949* (1.072)	-2.088* (1.137)
Survey Weights	NO	YES	NO	YES	YES	YES
Imputation	Listwise Deletion	Listwise Deletion	Simple	Simple Mean	Simple	Multiple
<i>N</i>	1686	1686	2081	2081	2081	2081
	1686	1686	2081	2081	2081	2081

Notes:

(i) Standard Errors given in parentheses

(ii) *, **, *** denote significance on a 10, 5, and 1% significance level respectively

3.2.6 Imputation

If the researcher is interested in making statements about the population at hand instead of the sample population, s/he needs to deal with the problem of item non-response. As with the problem of unit non-response, the problem of item non-response would not lead to biased estimates if it were a random process. Most of the time one can not assume that item non-response is a random process. There are many reasons why the process might very likely be non-random: Older people may not be able to answer questions relating to past events, more wealthy people tend to be unwilling to answer all questions related to their wealth, and/or the trust between certain interviewers and respondents might be lower than between others.

In order to control for the item non-response process (see Rubin and Little 1986 for details) in order to obtain unbiased estimates we will distinguish the four most common choices researchers make:

1. Ignore the item non-response process, also known as listwise deletion
2. Include the item non-response process in the model
3. Impute missing values first and use the imputed data
4. Use data already imputed by the data producers

While for a long time solution 1 was chosen by most researchers, solution 4 seems to be the most common one today. Nevertheless the quality of imputations and the methods used differ quite strongly over different surveys.

Solution 1 In general it is a bad idea to simply **ignore item non-response**. In most models, many variables are used and almost always the missing pattern in those variables is different for different observations. Nevertheless if only one variable is missing for a certain observation, the observation is excluded from the analysis. In most statistical software packages

(like STATA, PASW or R) this is the default procedure, and it implies that a huge amount of available information is discarded. Assume that you have a sample with 100 observations and 10 variables, ten missing patterns which are mutually exclusive and ten groups of households with a certain missing pattern for which always 1 variable is missing. That implies that out of the 1000 data points, 900 (90%) are non-missing - but at the same time, if you specify a model including all ten variables, you can not estimate anything because the number of usable observations is zero. Even if one is left with a reasonable number of observations, the estimates will be biased as soon as the item non-response process is non-random.

Solutions 2 and 3 To include the item non-response process in the model is (in theory) maybe the best approach because one has full control on how the item non-response process is modelled, and depending on the particular research question, the optimal solution to modelling this process might vary. The analyst can impute her/himself which is essentially the same, especially if a new imputation model is designed in order to fit a certain research question at hand. In both cases one can take care that all the relevant variables used later in the analysis are included in the imputation process and is therefore not dependent on the choice of the data producers which might not fit the research project at hand. But there are two straight forward problems researchers typically face when pursuing that approach: First, the available methods (such as the Heckman Selection Models) are grounded on strong assumptions, and second (and more severe), most research datasets do not include information which might be very relevant for the item non-response process. This information consists of information on the sampling, the interviewers, certain questions especially designed for the use in imputations as e.g. an estimate of monthly total net household income provided by the respondent, while all income sources are in general covered separately. Not being able to use this information - which might be available to the data producers and indeed used by them when imputing the missing values - might lead to a misguided and misspecified modelling of the item non-response process.

Solution 4 The main advantage of using already imputed data provided by the data producers is that they know the data much better than any analyst could and more importantly they are often able to use much more information which might be important for the item non-response process but is not available in the research dataset (information on sampling, interviewers, paradata, etc.). But the main problem is that they can not anticipate all possible research projects and therefore the imputation strategy used will fit well for some but not for all research purposes. There are very few surveys using state of the art chained equation multiple imputation approaches which ensure that (i) the joint distribution is taken into account properly, i.e. misspecified imputation does not destroy observed relationships between variables and (ii) that the uncertainty which comes with imputation can be taken into account properly when analysing the data, i.e. multiple imputations (instead of one) for every missing value reflect the uncertainty about the estimated (but not observed) value. The US Survey of Consumer Finances (SCF), the spanish Encuesta Financiera de las Familias (EFF) or the Survey of Health, Ageing and Retirement in Europe (SHARE) are some of the rare examples of surveys providing multiple imputations based on chained equations.

Why should the analyst care about imputation? Ignoring missing values (listwise deletion) leads to biased estimates as soon as the item non-response process is non-random. Furthermore, it discards a huge amount of observed information. Either modelling the item non-response process in the analysis or using imputed datasets therefore is preferable. The question therefore boils down to whether or not the researcher has enough information to model the item non-response process or impute on her/his own or should s/he rely on imputations done by the data producers?

The answer to this question depends heavily on three things:

- The quality of the available imputations
- The difference between the available information to the data producers (statistical

dataset) and to the researcher (research dataset)

- The research project at hand

Good imputations should take into account the joint distribution of the observed variables and use a broad conditioning approach to ensure the stability of relationships before and after imputation. They should also reflect the uncertainty of the imputed values, i.e. multiple instead of single imputations and should be produced in a consistent way, i.e. multiple imputation with chained equations instead of a mixture of different models to impute different variables independent of each other.

Good imputations need to take into account information on the sampling design as well as paradata. If such information is not available in the research dataset, it is hard to do better than data producers who had access to this type of information. The same is true for including the item non-response process in a model. Furthermore, the research project at hand is important. Let us assume that we want to measure the effect of education on wages, but for whatever stupid reason the data producers imputed wages using a simple regression model or hot-deck to impute but without taking education and its determinants into account. Then the relationship between education and wages can be severely damaged in the imputed dataset. Therefore it is always a good idea to look how data was imputed before using it.

Table 3.2 illustrates the impact of the use of listwise deletion, i.e. assuming the item non-response process to be random, and three different imputation-techniques, “simple mean” imputation, i.e. imputing missing values with means of non-missing observations, “simple” regression based equation (already using a chained equation model), and “multiple” imputation by chained equations provided by the survey producers (see Albacete 2012). Whereas the regressions based listwise deletion as well as simple mean imputation (see column OLS2 and column OLS4) lead to eight highly significant coefficients of pretty similar size, only 5 of them remain significant if more sophisticated imputation techniques are used (see column

OLS5 and OLS6). This shows that the standard errors of the former regressions - without proper imputation - are severely underestimated. In the case of the Inheritance dummy even the size of the coefficient changes significantly showing that the item-non response process was definitely non-random and therefore can lead to severe bias if it is not properly controlled for. Comparing the columns OLS5 and OLS6, which are both using up to date chained equation imputations, reveals that if the uncertainty which comes with imputed values is not controlled for, i.e. if simple instead of multiple imputation is used, the standard errors are still underestimated.

3.3 Final Remarks

In economics, researchers use many different types of data to test theoretical hypotheses or to describe some empirical facts. The goal is usually to say something about a certain target population. All of these data - if register data or survey data - is produced by humans and the way that it is produced has certain implications for the results the researcher will get.

It is therefore of utmost importance to be aware and if necessary take into account the data production process in the analysis. That includes the definition of the primary unit the data refers to (e.g. household sector, households, individuals, etc.), the general way of data gathering (e.g. register or survey data) and several other specific processes (e.g. interview mode in the case of survey data). Furthermore, especially in micro-data issues of coverage, unit- and item non-response play an important role. Therefore if possible one always needs to deal with questions of weighting and imputation in order to get a correct picture of a certain target-population. Very often one faces a trade off between bias and variance (e.g. trimming weights in data production or using weights in data analyses).

In the last two decades survey providers are increasingly taking care of those complex issues. While years ago only post-stratification weights and listwise deletion were common, today more

complex weighting and imputation techniques are used by a rising number of data providers. These tools allow for unbiased estimates (e.g. using proper multiple imputations or modelling the non response process) and make correct variance estimation feasible (e.g. using sampling information and linearization or replicate weights).

These methods make life a little more complicated for the researcher because using a multiple imputed dataset and taking care of the complex sampling structure of a survey implies some extra work. Even though most of the research-datasets available still do not allow to obtain a correct variance estimation (as they (i) do not include sampling information or provide replicate weights and (ii) use simple imputation methods which lead underestimated standard errors) the possibility to do so is rising fairly quickly.

We have tried to show that researchers should spend time in looking into and understanding the datasets they are working with, and to be aware of the possible impacts data production has on the results they are obtaining, which might be much larger than e.g. the choice of a slightly different econometric specification or model, an activity many researchers invest a large amount of time.

3.4 Literature

Albacete, N. 2012. Multiple Imputation in the Austrian Household Survey of Housing Wealth 2008. OeNB. Working Paper 167.

Angrist, Joshua and Pischke, Jörn-Steffen. 2009. Mostly harmless econometrics - an empiricist's companion. Princeton, NJ [u.a.]: Princeton Univ. Press.

Cameron, Collin A. and Pravin K. Trivedi. 2010. Microeconometrics Using Stata, Revised Edition. StataCorp LP, Stata Press books.

De Leeuw, E. D., J. J. Hox, and D. A. Dillman. 2008. International Handbook of Survey Methodology. European Association of Methodology.

Faiella, Ivan. 2010. The use of survey weights in regression analysis. Termini di discussione (Economic working papers). Bank of Italy, Economic Research Department.

Fessler, Pirmin, Maximilian Kasy and Peter Lindner. 2011. Identifying the causal effects of the interview mode on item non-response and income inequality measures in the Austrian EU-SILC survey – a quasi-experimental approach. Mimeo.

Groves et al. 2004. Survey Methodology. Wiley Series in Methodology. Wiley.

Hauser, Richard. 2007. Probleme des deutschen Beitrags zu EU-SILC aus der Sicht der Wissenschaft - Ein Vergleich von EU-SILC, Mikrozensus und SOEP, SOEP papers on Multi-disciplinary Panel Data Research, No. 69.

Kish, Leslie. 1995. Survey Sampling. Wiley Classics Library. Wiley.

Lohmann, Henning. 2011. Comparability of EU-SILC survey and register data: The relationship among employment, earnings and poverty. Journal of European Social Policy 2011, 21(1): 1-18.

Rubin, Donald B. and Roderick J. A. Little. 1986. Statistical Analysis with Missing Data. John Wiley & Sons, Inc. New York, NY, USA.

Chapter 4

A Note on the Effects of Higher Education on Earnings

4.1 Introduction

Over the last several decades a vast literature has been produced documenting the positive relationship between the education and earnings of individuals. Nevertheless, as social sciences suffer from the fact that they have no equivalent to laboratory experiments, it is difficult to make claims about a causal effect between the two. The fact that highly educated individuals earn more than individuals with lower levels of education is not only a result of education itself but the difference in earnings is at least partly due to selection bias. If more highly educated individuals earn more than lower educated individuals due to a stronger skill set, for example, the effect of education measured by the difference in observable earnings would overstate a causal effect. Many confounding factors might lead to selection bias overstating a causal effect. It has also been argued that ability bias might be small compared to other possible biases and that the causal effect would be understated by between-group earnings differences (Card (1999)).

Different econometric techniques have been used to control for selection bias in order to give estimates a causal interpretation. Most prominent in the literature on education and earnings are instrumental variable approaches based on either institutional features of the school system or family background, siblings and twin datasets as well as direct controls for abilities or family background (see Ashenfelter et al. (1999) and Card (1999) for extensive surveys of the literature.)

Card (1999) shows that family background (measured e.g. using parental education) is in general not a good instrument in that case, as likely a direct causal relationship between the instrument and the outcome exists and therefore the exclusion restriction (necessary for valid IV-estimation) does not hold. However, one can use parental education as a control to at least partly eliminate selection bias (Card 1999, page 1825).

One important issue when dealing with the causal effects of education on earnings and therefore with selection bias is the fact that the effect itself needs to be regarded as a random variable which might vary considerably with different socioeconomic and personal characteristics of individuals. Models differ extensively in the way they allow effects to be heterogenous. Blundell et al. (2005) provide an extensive review of the different workhorses in this strand of literature, which are OLS, IV, Control Functions and Matching.

The first of the few papers on the returns to education in Austria was Christl (1984), addressing differences in returns to education across occupations in a Mincerian framework. Boss et al. (1997) compare returns to education among private and public sector workers and Fersterer and Winter-Ebmer (1999) present estimates of the returns to education in Austria from 1981 to 1997. In their follow-up work Fersterer and Winter-Ebmer (2003) show that returns to education have been falling during this period.

In this paper we employ Austrian HSHW 2008 data in order to check if the returns to education in Austria as typically estimated in the literature might be overestimated due to the lack of control variables for social background. Our dataset includes parental education

which allows us to at least partly control for selection bias due to the social background of the individuals. None of the paper on returns to education for Austria covers such a control. In section 4.2 we discuss our empirical strategy. Section 4.3 describes the data we use. In section 4.4 we show and discuss results and in section 4.5 we conclude.

4.2 Empirical Strategy

4.2.1 Potential Outcome Approach

Most estimates of the returns to education literature use some type of the Mincerian setting

$$y_i = \alpha + \beta_1 S_i + \beta_2 E_i + \beta_3 E_i^2 + \varepsilon_i \quad (4.2.1)$$

where y_i are (log) earnings, S_i years of education and E_i years of working experience of individual i . If working experience is not available for some reason most researchers use the age of the individual minus S_i minus 6 as a proxy for E_i . If, as in many datasets, only discrete variables are available, dummies and interactions between them are used (see Card (1999)).

Sometimes industry or occupational dummies are also included in this type of earnings regression. However if one is interested in the (causal) effect of education on earnings (returns to education), that is a bad idea, as it adds additional bias. All variables which are themselves potential outcomes of the educational attainment are not proper controls (Angrist and Pischke (2009)).

Furthermore, most earnings regressions are plagued by two problems if one wants to give the resulting estimates of the returns to education a causal interpretation. First, educational attainment is not randomly distributed in the population and the most important controls for that selection are missing. Second, most models are not flexible enough and therefore the resulting estimates for the returns to education result from the strong linearity assumptions

and extrapolation outside the common support, even in terms of the included controls.

As shown in many empirical papers as (see Hertz et al. (2007) for an international comparison and Fessler et al. (2012) for results for Austria), parental education is highly correlated with educational attainment of the descendants. Many studies interested in the causal effect of education on earnings therefore use parental education to control for selection bias.

As selection bias in our empirical exercise results from the possibility that individuals with university degree might have had higher earnings anyway, (because selection to university is non-random) we use (i) parental education as control for selection bias and (ii) a flexible non-parametric model which allows us to restrict our analysis to the common support of university education with relation to our controls.

In this section we use a potential outcome approach as described in recent microeconomic literature on program evaluation and the measurement of causal effects in quasi-experimental designs, which can be found in a number of recent books (e.g. Morgan and Winship (2007), Angrist and Pischke (2009), Abbring and Heckman (2007)) and papers (e.g. Angrist and Krueger (1998), Blundell and Dias (2002) and Imbens and Wooldridge (2009)).

As we can not assume random assignment of university degrees, we need to control for selection into higher levels of education in order to justify a causal interpretation of estimated effects. Under the Conditional Independence Assumption (CIA)¹, treatment assignment is ignorable and independent of the potential outcomes,

$$P\{T_i|X_i, Y_{0i}, Y_{1i}\} = P(T_i|X_i) \tag{4.2.2}$$

and selection bias disappears. Under this assumption differences between the treatment

¹also referred to as unconfoundedness assumption and given by $\{Y_{0i}, Y_{1i}\} \perp\!\!\!\perp T_i|X_i$

and the control group have a causal interpretation, because educational attainment conditional on selection variables X_i is random. The remaining question is how to select the vector of covariates X_i . This set of controls needs to be chosen in a certain way to reasonably ensure the CIA but without introducing bad control bias, which is another form of selection bias (Angrist and Pischke 2009). In our case, good controls are variables which are themselves not a (possible) outcome of university degrees but are relevant for the assignment of the treatment as well as for the outcome. A safe way to prevent the danger of using controls which could induce selection bias is to use controls which are themselves already fixed at the time of obtaining a university degree or are strictly exogenous variables in relation to obtaining one (Imbens and Wooldridge 2009).

Note that in this paper we do not claim to estimate a pure causal effect as we are sure that there are other confounding factors we are not able to control for. However, what we can show is, that existing estimates of returns to education in Austria without proper control for social background are likely to be severely upward biased.

4.2.2 Average Treatment Effects on the Treated

We will focus on the average treatment effect on the treated (ATT), which is given by

$$ATT = \frac{1}{N_{T_1 \in \mathcal{C}}} \sum_{i \in \{T_i = 1 \cap \mathcal{C}\}} Y_{1i} - \hat{Y}_{0i}, \quad (4.2.3)$$

where $N_{T_1 \in \mathcal{C}}$ is the number of treated individuals in the common support \mathcal{C} and $\frac{1}{N_{T_1 \in \mathcal{C}}} \sum \hat{Y}_{0i}$ is the estimate of the potential outcome for the treated under non-treatment $\mathbb{E}[Y_{0i}|T_i = 1]$. The common support \mathcal{C} is defined by all covariate combinations (interactions), i.e. realizations of the covariate vector X ² where $0 < P(T_i = 1) < 1$. In other words, the common support is all resulting cells where both treated and non-treated individuals are observed (see Blundell

²Note that this strategy is only feasible if all covariate are discrete, or else the set of covariate cells is not finite.

et al. (2005)). Therefore realizations of X where $P(T_i = 1) = 0$ (no treated (no people with university degrees)) or $P(T_i = 1) = 1$ (all treated (only people with university degrees)) are excluded from the calculation of the ATT.

As long as X consists of discrete variables, the average treatment effect as in equation 4.2.3 can in principle be estimated by both OLS (saturated in X) and exact matching procedures (matched on X). The difference between estimating the ATT via OLS and via the matching process lies in the different weighting procedures the estimators use. While in the case of the matching estimator the covariate-distribution among the treated is used to combine covariate-specific effects into a single ATT, the regression function produces a variance weighted average of these effects. The matching estimator gives the most weight to the observations most likely to be treated whereas OLS gives the most weight to covariate combinations where the conditional variance of the treatment is maximized. The more the treatment effect varies across cells the more divergence there is in the resulting estimates.

In both the matching estimator and the OLS estimator using saturated controls X , only observations on the common support are used. In the matching estimator the observations in cells where the support of treated and non-treated is not overlapping cannot be matched and in OLS their weight is zero³.

As soon as continuous covariates are involved OLS saturated in X as well as exact matching is no longer feasible, as given a finite dataset the probability to find a match (with an identical realization of X) for a certain treated observations goes to zero. Several possibilities to control the necessary extrapolation outside the common support are feasible, for example via very flexible OLS models including higher order polynomials and as many interactions as possible or propensity score matching using a certain functional form to predict treatment

³The weight given by OLS to the covariate-specific treatment effect in a certain cell is $[P(X_i = x|T_i = 1)(1 - P(X_i = x|T_i = 1))]P(X_i = x)$, where x are certain realizations of X_i ; see Angrist and Pischke (2009).

probability and then match on the treatment probability.

In this empirical application we use coarsened exact matching (CEM) to estimate the ATT. Iacus et al. (2008) developed a method to temporarily coarsen data based on ex-ante user choice, match on this coarsened data, and then run the analysis on the uncoarsened data⁴. The advantage with relation to OLS or other matching methods is that it allows the user to define the common support \mathcal{C} via coarsening continuous variables by ex-ante user choice and then match the treated to their non-treated counterparts without using a parametric model. The analysis can then be conducted on the uncoarsened data which allows one to control for remaining imbalances in X due to coarsening of continuous variables.

4.3 The data

We use data from the Household Survey on Housing Wealth 2008 (HSHW 2008) of the Austrian Central Bank (Oesterreichische Nationalbank)⁵. It is an arguably representative household survey investigating the housing wealth of Austrian households. The respondents were either the owners or tenants of the respective household's real estate at the time of the interview. The survey focused on the ownership of the respective house/apartment and of additional real estate belonging to any of the household members as well as on the related liabilities owed by the household. Furthermore, detailed socio-economic characteristics, including the earnings (employed and self-employed monthly net-earnings) were compiled. The questionnaire contained a total of 168 questions, 28 of which were related to socio-economic characteristics (additionally, 8 questions had to be answered personally by the interviewers themselves). Finally, a response rate of 65,1% (2,081 observations) was accomplished.

The HSHW 2008 uses a stratified multistage cluster address random sample. It was car-

⁴We use *cem*, a STATA program provided by Matthew Blackwell, Stefano Iacus, Gary King and Giuseppe Porro, see <http://ideas.repec.org/c/boc/bocode/s457127.html>

⁵The HSHW 2008 fieldwork was conducted by the Institute for Empirical Social Studies (IFES).

ried out using a computer-assisted personal interviewing (CAPI) method, which allows for immediate plausibility checks during the course of the interview, thus making it possible to correct for inconsistencies right away. Multiple imputations using chained equations were used to impute missing values (see Wagner and Zottel (2009) for further general informations on the survey and specifically Albacete (2012) for informations on the multiple imputations).

The HSHW includes earnings, age, educational attainment as well as the educational attainment of both the mother and the father, for all respondents. We use the HSHW because it is one of the few datasets for Austria which includes information on parental education⁶.

However, the HSHW 2008 is focused on household wealth, and therefore not an optimal source for individual earnings. Nevertheless a question asking for the net monthly earnings from employed as well as self-employed (asking for an average over the last six months) of the respondent - which is the owner or tenant of the primary residence - is included. As we do not have hourly wages and the question with regard to working time per week is categorical we only use the sub-sample of individuals working 38.5 (Austrian full time) or more hours per week (as a result, 1018 out of 2081 observations are excluded). As university education is the main variable of interest in our analysis we drop individuals who are younger than 23, a reasonable lower age bound in order to have already been able to finish university education (dropping an additional 45 observations). We also drop individuals older than 65 years of age, the official retirement age (dropping another 6 observations). Finally we drop all individuals with monthly net earnings of less than 800 Euro to exclude outliers and reduce possible measurement error (additional 142 households are excluded). This leaves us 870 individuals for the analysis.

⁶EU-SILC 2005 contains the information on parental education for Austria. By examining the dataset, we found that there must have been some problems in the fieldwork concerning these variables, because their quality is insufficient for scientific analyses. Furthermore, the microcensus of 1996 contains educational background of parents (Fessler et al. (2012))

Table 4.1 shows a comparison of our main variables with the EU-SILC 2008 dataset. Experience is calculated in the standard ways by individuals age minus years of schooling minus six (see Card (1999)), where years of schooling are calculated from the discrete educational attainment variables in both surveys. EU-SILC, which is focused on income and therefore a more reliable source for individual incomes, collects yearly net earnings of all employed and self-employed individuals (aged 16 and older) living in the surveyed households. To get a proxy for monthly earnings we divide the yearly values by 12 in case of self-employed earnings and by 14 in case of employed earnings, which reflects the fact that employed individuals get a 13th and 14th extra salary⁷. Furthermore we used the same restrictions for individuals in the EU-SILC dataset as described for the HSHW, i.e. only full-time working individuals aged between 23 and 65 and with monthly net earnings of 800 Euro or more. That leaves us 3492 individuals for our analysis.

Table 4.1: Summary Statistics of Variables

	HSHW 2008				EU-SILC 2008			
	Min	Median	Mean	Max	Min	Median	Mean	Max
Earnings	800.00	1,700.00	2,068.28	10,976.00	800.00	1,612.07	1,878.80	72,319.88
Age	23.00	43.00	42.03	64.00	23.00	42.00	41.71	65.00
University	0.00	0.00	0.14	1.00	0.00	0.00	0.15	1.00
Experience	2.00	25.75	24.71	46.00	1.00	25.00	24.21	49.00
Female	0.00	0.00	0.40	1.00	0.00	0.00	0.30	1.00

Note that besides the differences in terms of measurement of the target variables, the samples do not refer to the exact same reference population. Both surveys are household surveys, but the HSHW 2008 only covers information on individual earnings for the respondent, whereas EU-SILC covers that information for all individuals aged 16 or more. Nevertheless as table 4.1 shows the variables seem to be surprisingly close to each other with regard to their distributions.

⁷Note that this leads to a slight upward bias, because fewer taxes have to be paid for the two extra salaries

With respect to earnings the EU-SILC includes some individuals with relatively high earnings, but only 5, who lie over the maximum of the HSHW. Also in terms of the share of individuals with university degree there is a slight difference of around 2 percentage points. All in all the educational attainments of respondents as well as those of their parents compare with other data sources from Statistics Austria quite well (see also Fessler and Schneebaum (2012)). The only larger difference in terms of our variables is the difference in the share of females. While the percentage of females is around 40% in the HSHW it is only around 30% in EU-SILC. In the Austrian Labour Force Survey 2008 the share women in total of full-time employed individuals is around 35% (Statistik Austria (2010)). The differences might be due to the aforementioned difference in the reference populations of the surveys.

4.4 Results

OLS Tables 4.2 and 4.3 show the results of three OLS regressions estimated with HSHW 2008 and EU-SILC 2008 data respectively. The first column *Diff* reports the unconditional effect of university degree on earnings. The second column *Mincerian I* refers to the classical Mincerian specification as given in equation 4.2.1 and the third column *Mincerian II* adds a female dummy to control for gender differences.

The (university) coefficient does not change significantly over the different specifications in both datasets. The difference in terms of the size between the EU-SILC and HSHW datasets is likely due to the differences in measurement and reference populations of both surveys.

However in both datasets the measured effects of university degree on earnings are positive, economically relevant and statistically significant in all specifications. A causal interpretation, however, is not in order as (i) the specifications rely on strong linearity assumptions and strong

Table 4.2: HSHW REGRESSIONS OF EARNINGS ON UNIVERSITY DEGREE

	Diff	Mincerian I	Mincerian II
University	0.211*** (0.044)	0.219*** (0.044)	0.219*** (0.044)
Experience		0.038*** (0.007)	0.038*** (0.007)
Experience squared		-0.001*** (0.000)	-0.001*** (0.000)
Female			-0.072** (0.032)
<i>N</i>	870	870	870

Notes: This table shows the average treatment effects of university degree on earnings resulting from OLS regressions using several sets of covariates as controls.

Table 4.3: SILC REGRESSIONS OF EARNINGS ON UNIVERSITY DEGREE

	Diff	Mincerian I	Mincerian II
University	0.350*** (0.022)	0.381*** (0.021)	0.396*** (0.021)
Experience		0.025*** (0.003)	0.024*** (0.003)
Experience squared		-0.000*** (0.000)	-0.000*** (0.000)
Female			-0.184*** (0.013)
<i>N</i>	3492	3492	3492

Notes: This table shows the average treatment effects of university degree on earnings resulting from OLS regressions using several sets of covariates as controls.

extrapolation outside the common support when averaging the treatment effect (regression coefficient), (ii) experience is already an outcome of our treatment variable university degree and therefore introduces bias, and (iii) selection bias is likely to be very severe as individuals with university degree might have higher earnings anyway and no proper controls for the selection to university are included.

CEM-ATT Table 4.5 shows the results of our coarsened exact matching estimates for ATT as given in equation 4.2.3. Our matching variables are female, age, experience and the maximum of parental education (fathers or mothers educational attainment measured in six categories). Female and parental education are discrete variables which we can match using exact matching. Experience and age are continuous variables which we coarse temporarily then match exactly on the coarsened dataset and analyse the dataset in its uncoarsened but matched form. In other words, we re-weight the dataset to balance the joint distribution of the covariates. Experience is temporarily coarsened to deciles of experience and age to categories (23-29; 30-39; 40-49; 50-59; 60-65). We use three different matching specifications which are shown in table 4.4. The first one (CEM I) is used as a benchmark case and does not control for parental education. In the second and third (CEM II and CEM III) we control for parental education in order to eliminate this confounding factor, which might (i) have a strong influence on the education level of the individuals and (ii) might have also a direct impact on the earnings of the individuals e.g. via network effects or abilities passed over generations but not reflected in education.

For the estimation of the ATT in column *CEM I* we only matched on experience and female to have a benchmark case to our OLS estimates. The estimate, although resulting from a more flexible model, is not different from the OLS estimates.

CEM II reports the estimated ATT after matching on experience, female and parental

Table 4.4: Coarsened Exact Matching

	CEM I		CEM II		CEM III	
Matching variables	Exp, Fem		Exp, Fem, PE		Age, Fem, PE	
No. of Covariate Combinations	16		82		50	
Common Support	15		54		36	
University	NO	YES	NO	YES	NO	YES
All	752	118	752	118	752	118
Matched	726	118	596	111	682	115
Unmatched	26	0	156	7	70	3

Table 4.5: Estimated ATT via coarsened exact matching

	CEM I	CEM II	CEM III
University	0.209*** (0.042)	0.129** (0.053)	0.126** (0.053)
Experience	0.036*** (0.008)	0.023** (0.011)	
Experience squared	-0.001*** (0.000)	-0.000 (0.000)	
Age			0.079*** (0.019)
Age squared			-0.001*** (0.000)
<i>N</i>	843	706	796

Notes: This table shows the average treatment effects of university degree on earnings resulting from coarsened exact matching estimators.

education. In this set-up individuals with university degree are only compared with their counterparts having no university degree but who also have - along with the same gender and experience - the exact same parental education. As parental education is a strong predictor for individual education it reduces selection bias up to a certain extent which leads to already much lower estimate of ATT.

CEM III reports the estimated ATT after matching on age, female and parental education. Since experience is already an outcome of education it is not a good control (see Angrist and Pischke 2009). Instead we use age itself. The resulting estimated ATT is again slightly lower.

In all three specifications we control for the remaining imbalances in the variables, which where temporarily coarsened, by including controls when calculating the ATTs.

Figures 4.1, 4.2 and 4.3 show the ATTs of *CEM I*, *CEM II* and *CEM III* respectively, but locally averaged over three different parental education categories. This allows us to look at the conditional effects (ATTs) underlying the estimated overall ATTs, which are given in table 4.5. Without control for parental education the effect for individuals with university degree is highest for those with high parental background (fig. 4.1, *CEM I*) and significant for individuals with low, medium and high parental education. With parental control the effects are on the contrary only significant for individuals with low and medium parental background. Furthermore when using age instead of experience effects are highest for those with low parental education (fig. 4.2 and 4.3, *CEM II* and *CEM III*).

All in all the effects without the use of parental control seem to be too large as well as misleading once looking at the underlying local effects. Once controlled for parental education we (i) find smaller returns to education and (ii) different heterogeneous effects which are basically turned around. The standard mincerian approach without control for parental education is based on underlying estimated higher returns for individuals with higher parental education.

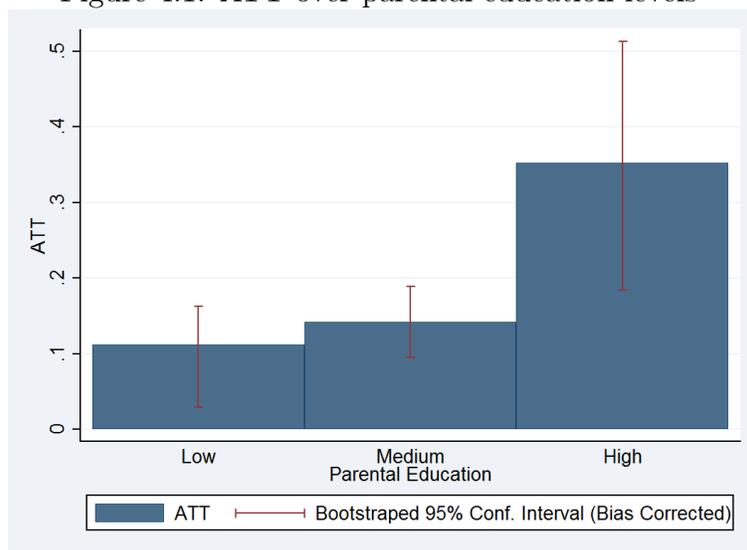
Once controlling for parental education we see that this result is mainly due to selection bias and the opposite is true: effects for individuals with high parental education are insignificant while for individuals with lower parental education we still find significant effects. This is due to the fact that the models without control for parental education compare apples to oranges. By averaging the treatment effects over the different groups of individuals with different parental backgrounds they ignore the fact that individuals with higher parental background tend to have higher earnings also because of other reasons than their own education (selection bias). If apples are only compared to apples and oranges to oranges, namely measuring the treatment effect of higher education only inside groups with the same parental background, and then aggregating that to an overall effect we find that the returns to education are indeed much smaller and reveal a different underlying structure of conditional effects.

For Austria, we therefore conclude that returns to education measured without control for parental educational background (or other social background variables), will typically overestimate the causal effect of educational attainment on earnings.

4.5 Conclusion

We employed HSHW 2008 data to illustrate that using standard mincerian equations when estimating so-called returns to education might overestimate the returns to education severely. Being able to control via parental education at least partly for selection bias due to social background we find indeed strong positive selection bias, i.e. descendants with highly educated parents (i) sort into higher education themselves more often (ii) have higher income already without university degree. Therefore without using parental education as a control, returns to education are likely overestimated. Furthermore we show that the effect of education on earnings might be very heterogenous over individuals with different social backgrounds. Once controlling for parental education we only find significant effects of education on income for individuals in lower parental education groups. For individuals with high parental education

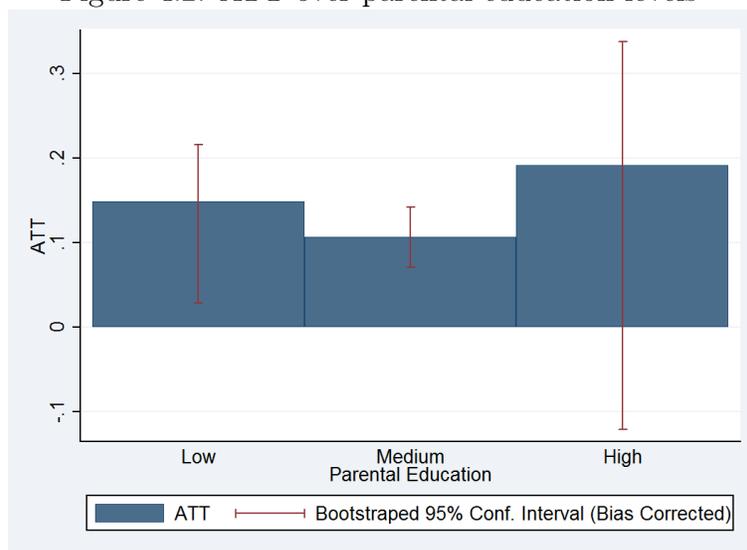
Figure 4.1: ATT over parental education levels



Notes:

(i) This figure shows the average treatment effects of university degree on earnings conditional on parental education resulting from CEM I. In Specification CEM I we control for (match on) experience and gender.

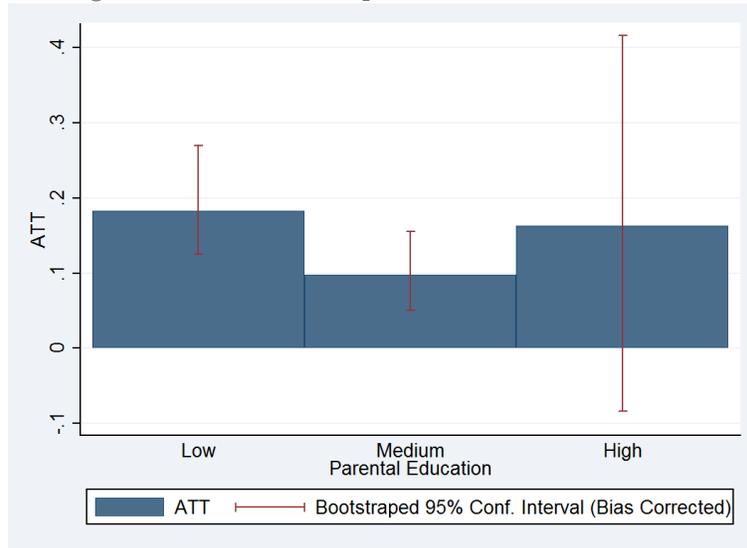
Figure 4.2: ATT over parental education levels



Notes:

(i) This figure shows the average treatment effects of university degree on earnings conditional on parental education resulting from CEM II. In Specification CEM II we control for (match on) experience, gender and parental education.

Figure 4.3: ATT over parental education levels



Notes:

(i) This figure shows the average treatment effects of university degree on earnings conditional on parental education resulting from CEM III. In Specification CEM III we control for (match on) age, gender and parental education.

a university degree does not have a significant effect on income, as there are obviously other forces at play that increase income without additional education.

We conclude that “returns to education” estimated by standard mincerian equations - omitting social background variables - are often biased. By controlling for parental education we show for Austria that for a large degree, the so-called “return” measured is actually due to selection bias. However we do not claim that the effect measured in this paper should be interpreted as a causal one as we can definitely not eliminate all selection bias. More research using (natural) (quasi-) experiments is needed to tackle this issue in Austria.

4.6 Literature

Abbring, J. H. and J. J. Heckman. 2007. Econometric Evaluation of Social Programs, Part III: Distributional Treatment Effects, Dynamic Treatment Effects, Dynamic Discrete Choice, and General Equilibrium Policy Evaluation, in: Handbook of Econometrics, ed. by J. Heckman, and E. Leamer, vol. 6B, chap. 72. Elsevier, 1 edn.

Albacete, N. 2012. Multiple Imputation in the Austrian Household Survey of Housing Wealth 2008. OeNB. Working Paper 167.

Angrist, J. and A. B. Krueger. 1998. Empirical Strategies in Labor Economics, Working papers 98-7, Massachusetts Institute of Technology (MIT), Department of Economics.

Angrist, J. and J.-S. Pischke. 2009. Mostly harmless econometrics. Princeton, NJ [u.a.]: Princeton Univ. Press.

Ashenfelter, O., C. Harmon and H. Oosterbeek. 1999. A review of estimates of the schooling/earnings relationship, with tests for publication bias, Labour Economics, 6(4), 453-470.

Boss, M., H. Hofer, P. Mitter and J. Zweimüller. 1997. Lebens Einkommen im Privaten und Öffentlichen Sektor. Institute for Advanced Studies, Vienna.

Blundell, R., L. Dearden and B. Sianesi. 2005. Evaluating the effect of education on earnings: models, methods and results from the National Child Development Survey, Journal Of The Royal Statistical Society Series A, 168(3), 473-512.

Blundell, R. and M. C. Dias. 2002. Alternative approaches to evaluation in empirical microeconomics, CeMMAP working papers CWP10/02, Centre for Microdata Methods and Practice, Institute for Fiscal Studies.

Card, D. 1999. The causal effect of education on earnings, in Handbook of Labor Economics, ed. by O. Ashenfelter, and D. Card, vol. 3 of Handbook of Labor Economics, chap. 30, pp. 1801-1863. Elsevier.

Christl, Josef. 1984. The explanatory power of the human capital earnings function. *Empirica* 11(1), 47-57.

Fersterer, Josef and R. Winter-Ebmer. 1999. Returns to Education: Evidence for Austria, Arbeitspapier No. 9913, University of Linz, Austria.

Fersterer, J. and R. Winter-Ebmer. 2003. Are Austrian returns to education falling over time?, *Labour Economics*, 10(1), 73-89.

Fessler, Pirmin, Peter Mooslechner, and Martin Schürz. 2012. Intergenerational Transmission of Educational Attainment in Austria, *Empirica: Volume 39, Issue 1 (2012)*, Page 65-86.

Fessler, P. and Schneebaum, A. 2012. Gender and Educational Attainment Across Generations in Austria, *Feminist Economics*, 18:1, 161-188.

Hertz, Tom; Tamara Jayasundera, Patrizio Piraino, Sibel Selcuk, Nicole Smith, and Alina Verashchagina. (2007). The Inheritance of Educational Inequality: International Comparisons and Fifty-Year Trends, *The B.E. Journal of Economic Analysis and Policy* 7(2): Article 10.

Iacus, S., G. King and G. Porro. 2008. Matching for causal inference without balance checking, Working paper series, Harvard.

Imbens, G., and J. Wooldridge. 2009. Recent Developments in the Econometrics of Program Evaluation., Journal of Economic Literature 47(1): 586.

Morgan, S., and C. Winship. 2007. Counterfactuals and Causal Inference. Cambridge, Cambridge University Press.

Statistik Austria. 2010. Arbeitsmarktstatistik Jahresergebnisse 2010, Mikrozensus- Arbeitskräfteerhebung, Schnellbericht 5.8.

Wagner, K. and S. Zottel. 2009. OeNB-Immobilienvermögenserhebung 2008 der privaten Haushalte Beschreibung, verwendete Methoden und Bewertungsansätze, Statistiken Q4/09, 45-65, Oesterreichische Nationalbank.

Appendix A

Abstracts

A.1 English Abstract

A.1.1 The Gendered Aspects of the Intergenerational Persistence of Educational Attainment in Austria

In many societies, childrens' education levels are heavily dependent on their parents' education, but that result can differ by the gender of the child. Using a Markovian approach, along with uni- and multivariate econometric techniques, we employ the Austrian Household Survey on Housing Wealth to show strong persistence in educational attainment that differs by the gender of the parent and the child. We find that the size of educational persistence varies over time in Austria and that the relevance of one's father's education is generally higher than that of the mother, once controlling for distributional differences. Further, the relationship between parents and children of the same gender is stronger than the cross-gender parent/child relationship. The educational penalty for females has been shrinking over time while educational mobility for both genders has increased over time.

A.1.2 Assets and Liabilities of Austrian Households

On macro level data one can see that assets as well as liabilities of private households increased rapidly in recent years. This is true for many OECD countries including Austria. The parallel movement on macro level of both is for itself not puzzling. Though to assess financial stability issues its of utmost importance to know which households are responsible for the growth of liabilities and assets. Are these different ones or the same ones? We try to asses this question and to gain deeper insights in what determines the participation in debt market and what determines the volume of debt. Conducting descriptive statistics and a Heckman Selection model on the basis of Survey data, we find that higher income, higher educated and households of larger size tend to hold debt more often and higher amounts of debt. Surprisingly the amount of debt is also rising with financial assets of a household.

A.1.3 The Role of Data Production in Survey Analysis

A large amount of economic research is based on analysis of observed data; this paper discusses some aspects of the relationship and disconnect between data collection and data analysis. It is usually the case that the data collectors are not the same individuals as those analysing the data, and in nearly all cases, collected data is heavily influenced by the individuals gathering the data. We show how the data production process can influence data analysis and its results, discuss the caveats of survey data analysis with respect to sampling, the survey mode, interviewing, paradata, weights and imputation and document how ignoring those issues might lead to severe bias as well as a misleading precision of estimates when analysing the data.

A.1.4 A Note on the Effects of Higher Education on Earnings

A vast literature exists estimating the effects of education on individual earnings. As those attaining a certain education most likely would earn also more without one than those without a certain education many approaches have been used to control for unobserved heterogeneity and resulting selection bias in order to give the estimated coefficients a causal interpretation.

Most prominent are instrumental variable approaches based on either institutional features of the school system or family background, siblings and twin datasets as well as direct controls for abilities or family background. In this paper we employ Austrian HSHW 2008 data which allows us to control for parental education in order to estimate the effect of university education on earnings. By allowing for heterogenous returns using non-parametric models we show that the effects of a university degree on earnings decreases sharply with model flexibility as well as with parental education. Using coarsened exact matching to establish common support and construct a meaningful counterfactual we find that the conditional average treatment effect of university education on earnings is significant given low parental education but insignificant given high parental education. This points towards severe bias with regard to so called returns to education in standard mincerian approaches which tend to ignore selection bias related to parental education.

A.2 German Abstract

A.2.1 The Gendered Aspects of the Intergenerational Persistence of Educational Attainment in Austria

Der Ausbildungsgrad von Eltern und Kindern ist in vielen Gesellschaften stark korreliert. Dieser Zusammenhang unterscheidet sich auf der Ebene der Geschlechter von Eltern und Kindern sowie deren Kombination. Auf Basis des österreichischen Household Survey on Housing Wealth zeigen wir dass (i) es eine starke Persistenz in Bezug auf die Bildung über die Generationen hinweg gibt (ii) die Persistenz über die zeigt stark variiert und (iii) starke Muster je nach Geschlecht der Eltern und Kinder aufweist. Im Allgemeinen ist die Korrelation zwischen Vätern und Kindern stärker als zwischen Müttern und Kinder. Weiters ist jene zwischen Vätern und Söhnen stärker als jene zwischen Vätern und Töchtern sowie jene zwischen Müttern und Töchtern stärker als jene zwischen Müttern und Söhnen. Trotz eines Schließens der Bildungslücke der Geschlechter bleiben diese Muster stabil.

A.2.2 Assets and Liabilities of Austrian Households

Auf Makroebene kann in den letzten Jahrzehnten ein starker Anstieg der Schulden, aber auch des Geldvermögens, privater Haushalte beobachtet werden. Doch handelt es sich um die Selben oder Andere Haushalte, die für diese Phänomene verantwortlich zeichnen? Diese Frage ist von enormer Bedeutung für die Finanzmarktstabilität. Wir versuchen auf Basis von Mikrodaten einen Einblick in die Determinanten der Verschuldung und deren Höhe zu erlangen. Mit Hilfe eines Heckman Selection Models zeigt sich, dass abgesehen vom Alter vor allem Einkommen, Bildung und Haushaltsgröße wichtige Determinanten für die Verschuldung darstellen. Bemerkenswerterweise steigt die Verschuldung aber auch mit dem Geldvermögen der Haushalte.

A.2.3 The Role of Data Production in Survey Analysis

Ein großer Bereich der empirischen Wirtschaftsforschung basiert auf der Analyse von Erhebungsdaten (Surveydaten). In diesem Papier besprechen wir die zentralen Aspekte des Verhältnisses und der Trennung zwischen Datenerfassung und Datenanalyse. Im Allgemeinen sind die Datenproduzenten nicht gleichzeitig jene Personen die die Daten analysieren. Wir besprechen, wie der Datenproduktionsprozess die Datenanalyse und ihre Resultate beeinflussen kann. Dabei werden die Themengebiete Stichprobenziehung, Erhebungsinstrumente, Paradata, Gewichte und Imputationen besprochen und gezeigt wie das Ignorieren derselben zu verzerrten Schätzern und irreführenden - meist zu kleinen - Standardfehlern derselben führen kann.

A.2.4 A Note on the Effects of Higher Education on Earnings

Eine umfangreiche ökonomische Literatur beschäftigt sich mit den Effekten der Ausbildung eines Individuums auf dessen Einkommen. Da jene Individuen, die ein bestimmtes Ausbildungsniveau erreichen wahrscheinlich auch mehr verdienen würden ohne diese Ausbildung als jene, die dieses Niveau gar nicht erreichen würden oder wollen, ist es notwendig für diesen Selektions-Bias zu kontrollieren. Wir zeigen, dass die in der "Returns to Education" Literatur besonders verbreiteten Mincer-Spezifikation tendenziell den Effekt der Ausbildung auf das Einkommen überschätzen. Anhand von flexiblen nicht-parametrischen Methoden kann zudem gezeigt werden dass der Effekt der Ausbildung sowohl mit der Flexibilität der Spezifikation als auch mit der Elternbildung abnimmt. Neben der Tatsache, dass die Standard Mincer Spezifikation den Selektions-Bias in Bezug auf die Elternbildung zu einer deutlichen Überschätzung des Effekts führt dokumentieren wir damit, dass Ausbildung vor allem für Personen mit niedriger Elternbildung zu höherem Einkommen führen kann während jene mit hoher Elternbildung auch ohne eigene Ausbildung schon zu einem relativ hohen Einkommen gelangen und daher der Bildungseffekt auf das Einkommen für diese Gruppe vernachlässigbar ist.

Appendix B

Curriculum Vitae

1160 Vienna

Austria

Date of Birth: October 06, 1980

Citizenship: Austria

Office: (+43) (1) 404 20 -7414

Private: (+43) (0) 6508243920

Office E-mail: pirmin.fessler@oenb.at

Private E-mail: pirminfessler@gmail.com

Education

- Ph.D. Program in Economics, University of Vienna, *expected to be finished in 2012*.
 - *Area of Dissertation*: Microeconometrics, Survey Data.
- M.A. Economics, University of Vienna, 2006.
 - *Master Thesis*: Scale effects in R&D based endogenous growth models.

Work Experience

- Economist, Monetary Unit, Economic Analysis Division, Economic Analysis and Research Department, Oesterreichische Nationalbank (Austrian Central Bank), 2007-Present.

- Research Assistant, Economic Analysis and Research Department, Oesterreichische Nationalbank (Austrian Central Bank), 2006-2007.

Academic Experience

- Lecturer, Seminars in Fundamentals of Economics, Economic Policy, Monetary Macroeconomics, University of Applied Sciences BFI Vienna (Fachhochschule des bfi Wien), 2010-Present.
- Research Assistant, Vienna University of Economics and Business Administration, Department for International Economics and Development, SFB-International Tax Coordination. June 2006-October 2006.
- Teaching Assistant, Vienna University of Technology, Institute for Mathematical Methods in Economics, Introduction to Economics, Winter Term 2005/2006.

Publications

Books

- Fessler, P., Hinsch S. (2011): *Wie funktioniert Wirtschaft? Eine kritische Einführung*, Promedia Verlag, Wien.
- Fessler, P. (2010): *Wachstumstheorie: Skaleneffekte in endogenen forschungs- und entwicklungs-basierten Wachstumsmodellen*, VDM Verlag Dr. Müller.

Contributions to Books

- Fessler, P., Mooslechner P., Wagner K., Schürz M. (2009): *Housing and financial wealth in Austria: What can survey data tell us for the analysis of financial stability issues?*, In:

Housing Market Challenges in Europe and the United States - Any solutions available?, Palgrave MacMillan.

Articles

- Fessler, P and Schneebaum, A. (2012): Gender and Educational Attainment Across Generations in Austria, *Feminist Economics*, 18:1, 161-188.
- Fessler, P., Mooslechner P., Schürz M. (2012): Intergenerational Transmission of Educational Attainment in Austria, *Empirica: Volume 39, Issue 1 (2012)*, Page 65-86.
- Andreasch, M., Fessler, P., Mooslechner, P., Schürz, M. (2011): Einkommen privater Haushalte auf Mikro- und Makroebene – HFCS versus VGR, *Statistiken Q2/11*.
- Fessler, P., Mooslechner P., Schürz M. (2011): Zur Messung des Pensionsvermögens im HFCS, *Statistiken, Q1/11*, p. 37-47. Oesterreichische Nationalbank (Austrian Central Bank).
- Fessler, P., Mooslechner P., Schürz M. (2010): Repräsentativität der ersten Welle des HFCS in Österreich, *Statistiken, Q4/10*, p. 48-62. Oesterreichische Nationalbank (Austrian Central Bank).
- Fessler, P., Mooslechner P., Schürz M. (2010): Zur Konzeption des Vermögens in der Erhebung des Eurosystems zu Finanzen und Konsum der privaten Haushalte, *Statistiken, Q3/10*, p. 40-55. Oesterreichische Nationalbank (Austrian Central Bank).
- Fessler, P. and Schürz M. (2010): Informationen zum kleinen Häuselbauer, *Wirtschaft und Gesellschaft, Heft 02/10*. Arbeiterkammer Wien.
- Fessler, P., Mooslechner P., Schürz M. (2010): Real Estate Inheritance in Austria, *Monetary Policy and the Economy, Q2/10*, p. 33-53. Oesterreichische Nationalbank (Austrian Central Bank).

- Albacete, N., Fessler, P. (2010): Stress Testing Austrian Households, Financial Market Stability Report 19, p.72-91. Oesterreichische Nationalbank (Austrian Central Bank).
- Andreasch, M., Fessler, P., Schürz M. (2009): Austrian Households Equity Capital - Evidence from Microdata, Monetary Policy and the Economy, Q4/09, p.61-78. Oesterreichische Nationalbank (Austrian Central Bank).
- Fessler, P., Mooslechner P., Schürz M. (2009): Interviewtechniken bei Erhebungen zu den Finanzen privater Haushalte im Euroraum, Statistiken, Q4/09, p.66-77. Oesterreichische Nationalbank (Austrian Central Bank).
- Fessler, P., Mooslechner P., Wagner K., Schürz M. (2009): Housing Wealth of Austrian Households, Monetary Policy and the Economy, Q2/09, p.104-126. Oesterreichische Nationalbank (Austrian Central Bank).
- Fessler, P., Mooslechner P., Schürz M. (2009): Stichprobenziehung bei Erhebungen zu den Finanzen privater Haushalte im Euroraum, Statistiken, Q2/09, p.51-62. Oesterreichische Nationalbank (Austrian Central Bank).
- Fessler, P., Mooslechner P., Schürz M. (2009): Herausforderungen der Forschung zur Finanzsituation privater Haushalte im Euroraum, Statistiken, Q1/09, p.57-66. Oesterreichische Nationalbank (Austrian Central Bank).
- Fenz, G. and Fessler, P. (2008): Wealth effects on Austrian Consumption, Monetary Policy and the Economy, Q4/08, p.68-84. Oesterreichische Nationalbank (Austrian Central Bank).
- Fessler, P. and M. Schürz. (2008): Stockholdings in Austria, Monetary Policy and the Economy, Q2/08, p.83-100. Oesterreichische Nationalbank (Austrian Central Bank).
- Fessler, P. and P. Mooslechner. (2008): Arme Schuldner - Reiche Schuldner? Haushaltsverschuldung und Geldvermögen privater Haushalte auf Basis von Mikrodaten. In: Intervention, European Journal of Economics and Economic Policies, 5(1), 31-45.

- Fessler, P., Schürz M., Wagner K., Weber B. (2007): Financial Capability of Austrian Households, Monetary Policy and the Economy Q3/07, 50-67, Oesterreichische Nationalbank (Austrian Central Bank).

Working Papers

- Fessler, P., Mooslechner P., Schürz M. (2008): How Inheritances Relate to Wealth Distribution? Theoretical Reasoning and Empirical Evidence on the Basis of LWS Data, Luxembourg Wealth Study Working Paper Series, Working Paper No. 6
- Fessler, P. (2006): Home Country Effects of Offshoring - A Critical Survey on Empirical Literature, Discussion Papers SFB International Tax Coordination, Nr. 23, December 14, 2006 SFB International Tax Coordination, University of Economics and Business Administration, Vienna

Invited Talks

Conference Presentations

- CIRET (Centre for International Research on Economic Tendency Surveys), New York, United States October 13-16, 2010, Stress Testing Austrian Households, with Nicolas Albacete.
- Brussels Representative Office (OeNB), September 7, 2010, Austrian Households Equity Capital, with Michael Andreasch.
- IARIW (The International Association for Research in Income and Wealth) 31th General Conference, St. Gallen, Switzerland August 23-30, 2010, with Martin Schürz.
- ECNIEQ (Society for the Study of Economic Inequality), Third meeting in Buenos Aires, July 20-23, 2009, Intergenerational Transmission of Educational Attainment in Austria, with Martin Schuerz.

- IAFFE (International Association for Feminist Economics), General Boston, 26-28 June 2009, Gender aspects in the intergenerational persistence of education, with Alyssa Schneebaum.
- Saving in Austria - Too little and too late?, Oesterreichische Nationalbank, 24 October 2008, Vienna, Savings in Austria with Clemens Jobst
- IARIW (The International Association for Research in Income and Wealth) 30th General Conference, Portoroz, Slovenia August 24-30, 2008, How Inheritances Relate to Wealth Distribution?, with Martin Schürz.
- Banca d'Italia, Enhancing Comparative Research on Household Finance, Roma, Italy, 5-7 July 2007, How Inheritances Relate to Wealth Distribution?, with Peter Mooslechner and Martin Schürz.

Seminars/Workshops

- Vienna University of Economics and Business Administration, Institute for Money and Finance, Topic: Wealth Inequality in Austria, June 2008.
- ECB (European Central Bank), September 13-14, 2010, Topic: Multiple Imputation in the HFCS (Household Finance and Consumption Survey).

Refereeing

- The Review of Income and Wealth

Further Professional Activities

Service

- Member of the Household Finance and Consumption Network, ECB (European Central Bank), Eurosystem, May 2008–Present.
- Member of the Advisory Board of LIS (Luxembourg Income Study), Luxembourg & New York, July 2011–Present.

Professional Memberships

- The International Association for Research in Income and Wealth, March 2008–Present.
- Society for the Study of Economic Inequality, June 2009–Present.

Summer/Winter Schools Attended

- Summer School, Microeconometrics, OENB (Oesterreichische Nationalbank), by Collin Cameron, August 2007.
- The fourth Winter School on Inequality and Social Welfare, Alba die Canazei, January 2009
- LIS (Luxembourg Income Study), Summer Workshop 2008, July 2008.
- Summer School, Survey Data Analysis, OENB (Oesterreichische Nationalbank), by Franco Peracci, August 2007.

Miscellaneous

Computer Skills

Linux, Mathematica, STATA, R, SPSS, L^AT_EX.

Special Interests

Philosophy, Politics, Sociology, Travelling (especially Latin America), Open Source Software.

Language Skills

German (native language), English (speak fluently and read/write with high proficiency), Spanish and French (speak and read with basic competence).

Appendix C

Acknowledgements

I thank my supervisors Robert M. Kunst and Christine Zulehner for their valuable comments and discussion. Furthermore several people contributed in one way or another - mainly discussions about the topics at hand or related topics - to this work. I am grateful to Nicolas Albacete, Verena Halsmayer, Max Kasy, Markus Knell, Peter Lindner, Ina Matt, Peter Mooslechner, Marlene Penz, Alyssa Schneebaum and the members of the ECB “Household Finance and Consumption Network”. In addition I want to thank two people who were especially helpful, inspiring and motivating and are outstanding - especially in the field of economics - because of their true commitment to the social science and their aim to increase and impart knowledge but at the same time being aware of its limits: Martin Schürz and Arthur Kennickell.