Inequality of Educational Outcome and Inequality of Educational Opportunity in the Netherlands during the 20th Century

Maarten L. Buis

Thesis committee:

prof.dr. R. Breen (Yale University)
dr.ir. C.H. Elzinga (VU University Amsterdam)
prof.dr. P.M. de Graaf (Tilburg University)
prof.dr. S.P. Jenkins (University of Essex)
prof.dr. T. Tam (Chinese University of Hong Kong)

ISBN: 978-90-8659472-6

VRIJE UNIVERSITEIT

Inequality of educational outcome and inequality of educational opportunity in the Netherlands during the 20th century

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad Doctor aan
de Vrije Universiteit Amsterdam,
op gezag van de rector magnificus
prof.dr. L.M. Bouter,
in het openbaar te verdedigen
ten overstaan van de promotiecommissie
van de faculteit der Sociale Wetenschappen
op vrijdag 2 juli 2010 om 13.45 uur
in de aula van de universiteit,
De Boelelaan 1105

door

Maarten Leendert Buis
geboren te Amsterdam

promotor: prof.dr. H.B.G. Ganzeboom
copromotor: dr. A.W. Hoogendoorn

Acknowledgments

While writing this dissertation I was fortunate to receive support in a variety of forms from many people, and I would like to use this opportunity to thank them. In particular I would like to thank my promotor Harry Ganzeboom for introducing me into the field of social stratification research and for his critical and constructive comments on my texts and analyses. In addition I thank my copromotor Adriaan Hoogendoorn for the frequent brain storm sessions (and the coffee).

Very important was the camaraderie and support from my former colleagues in the department, especially of those with whom I shared an office (or whose office I visited so often that that was *de facto* the case): Kadri Täht, Tina Glasner, Yfke Ongena, and Irma Mooi-Reçi. They were for a large part responsible for the day to day fun.

My friends helped maintain my sense of perspective. In particular, Anne Potters who contributed her down to earth insights into what life is like outside academia, and who mobilized the "language nerds" from her circle of friends, to help me, a "numbers nerd", with the difficult task of writing. The people of woongroep Pantha Rhei were of great practical help, as they made sure I always had a warm meal, regardless of how late I got home.

Finally, I thank my mother, father, and sister for their support and confidence in me.

Contents

Chapter 1

Introduction

It is common practice to start studies on education with a claim that education is an important determinant of later life chances (for example Mare, 1981; Shavit and Blossfeld, 1993). Instead of repeating this claim I will report the following official statistics for the Netherlands: the unemployment rate in 2006 for persons with only primary education is 12.2% versus 3.7% for persons with a university degree; in 1998 29% of women aged between 34 and 38 with a university degree expected to remain childless versus 16% for women with primary or lower secondary education; men who were born in 2008 are expected to live 50.2 years in good health if they only complete primary education versus 69.0 years if they attain a degree in tertiary education (Statistics Netherlands, 2008). These statistics sufficiently illustrate the importance of education for a wide range of domains in a person's life and the role of education as the primary stratification mechanism in modern societies.

If a resource like education is this important, then the distribution of this resource is certainly worth studying. There is a long list of literature that has done just that, and it shows that educational attainment is unequally distributed among persons with different family backgrounds, in particular that persons from more privileged families tend to obtain more education than persons from less privileged backgrounds (Hout and DiPrete, 2006; Breen and Jonsson, 2005). In this dissertation I will try to contribute to the study of this inequality in access to education. I will focus on two types of inequality of access to education and the relationship between these types. The first type of inequality in access to education is the inequality as it arises during the process of attaining education. This is usually captured by studying the effect of family background on the probabilities of passing from one level of education to the next, and I will call this Inequality of Educational Opportunity or IEOpp[1] The second type of inequality in terms of access to education is the inequality in the end result of the educational selection process. This is usually captured by studying the effect of family background on the highest achieved level of education, and I will call this Inequality of Educational Outcome or IEOut. The dominant issue in this literature is whether or not the IEOpp and IEOut have changed over time, and in particular, whether they

[1]The term Inequality of Educational Opportunity (usually abbreviated to IEO) was already used by Boudon (1974) and Mare (1981), where it is used as a more generic term for inequality of access to education. However, in the studies by Boudon (1974) and Mare (1981) the effects of family background on the probabilities of passing from one level to next are claimed to be a more "pure" representation of IEO.

have decreased over time. A common finding for the Netherlands has been that for this country there has been a gradual and long-term decline of inequality in both IEOpp and IEOut during the course of the 20^{th} century (De Graaf and Ganzeboom, 1993; Ganzeboom and Luijkx, 2004b). These results have been obtained using a continually extending database of pooled cross-section data, most recently consisting of over 50 surveys held in the Netherlands since 1958 covering cohorts born throughout almost the entire 20^{th} century. The aim of the studies collected in this dissertation is to re-assess and extend the evidence in these earlier studies, primarily from methodological points of view. Overall, the research question guiding the separate studies in this dissertation can be formulated as follows:

> To what extent, how, and when has a trend toward less inequality in educational opportunities and in educational outcomes of persons from different family backgrounds occurred in the Netherlands?

The first step undertaken to elaborate and answer this general research question is to provide an overview of the trends in IEOpp and IEOut following the protocol used in an influential international comparative project headed by Shavit and Blossfeld (1993), but using the most recent data available on the Netherlands. This analysis will be a replication of the Dutch contribution to this project by De Graaf and Ganzeboom (1993). Such a replication is useful in its own right, but will also function as the point of departure to which all results in the subsequent chapters can be compared. This replication will be presented in Chapter 2.

The subsequent chapters in this dissertation will each discuss a way of improving this ‘default’ method and the consequences of these methodological innovations for the estimated trends. Chapters 3, 4, and 5 discuss various ways of improving the estimates of IEOut. Chapter 3 will introduce a way of improving the scale on which the highest achieved level of education is measured. Chapter 4 will focus on how best to measure any changes in the trend in IEOut. Chapter 5 will investigate the relative influences of different indicators of family socioeconomic status. Chapter 6 will introduce a way to integrate the analysis of IEOpp and IEOut, thus allowing one to make the best use of the complementary nature of these two representations of inequality in access to education. This integration will also provide a substantive interpretation of the effect of educational expansion — the fact that people from more recent cohorts attain, on average, higher levels of education than people from older cohorts — on IEOut. Finally, Chapter 7 will propose a way of dealing with an influential critique by Cameron and Heckman (1998) on the most common method of estimating IEOpps.

Chapter 3 will focus on the scaling of education. In order to study IEOut — that is, the effect of family background on the highest achieved level of education

— one needs to assign values to each level of education[2]. In Chapter 3 these values will be empirically estimated such that education optimally predicts the respondent's occupational status. Most previous studies of IEOut use an *a priori* scale of education that is loosely based on the number of years it should take a 'standard' student to finish that level. Such a scale conflates two related but distinct concepts: the duration and the value of education. Another issue is that such an *a priori* scale assumes that the values are constant over time, while there is an influential hypothesis that states that the value of the higher educational categories have declined, so-called diploma inflation. This hypothesis is based on the fact that people born more recently on average achieve much higher levels of education than people born longer ago. As a consequence the number of higher educated persons has increased, which has led to the prediction that the value of their education has declined. Chapter 3 will test whether the estimated values of the levels of education have actually changed over time, and compare the estimated values with commonly used *a priori* values.

Chapter 4 will focus on the question of whether or not the trend in the effect of family background on educational outcomes has changed over time. Existing literature has occasionally tested for the presence of curvilinear (accelerated of decelerated) trends, but found little or no supporting evidence (De Graaf and Ganzeboom, 1990; De Graaf and Luijkx, 1992; De Graaf and Ganzeboom, 1993; De Graaf and Luijkx, 1995; Ganzeboom, 1996). This is implausible: if the long-term trend is towards lower association between social background and educational achievement, one would expect a slow-down of this trend at some point, as otherwise a continuing linear trend would lead to a negative association between social background and educational achievement. In Chapter 4 I examine whether such a non-linear development has already occurred, using local regression models that appear to be new to this field.

Chapter 5 will focus on the relative importance of different types of family background, in particular, the education and occupational status of both parents. It is probable that the relative contributions of these resources have changed over time. Two such changes are expected from the literature: First, economic resources (parental occupational status) are predicted to have become less important relative to cultural resources (parental education). The effect of economic resources are expected to decline, because the combination of economic growth and an increase in government subsidies is likely to have decreased the negative influence of poverty on attaining education. A similar decline in effect of the cultural resources is not expected, leading to the expectation of a increase of importance of the cultural resources relative to economic resources (De Graaf and Ganzeboom, 1993). Second, the resources contributed by the mother are likely to have increased in importance relative to the resources con-

[2] In studies of IEOpp, a similar issue arises with respect to the rank order of the transitions analysed, but this presents less of a puzzle as this order is usually institutionally determined.

tributed by the father due to the changing roles of men and women in society (Korupp et al., 2002). These hypotheses are of substantive interest in their own right, but they also have an important practical consequence for social stratification research. Studies in this field often use only one of these resources, most typically father's occupational status, as in indicator of family socioeconomic status. If the relative contributions of the different resources have changed over time, then trends in IEOpp or IEOut found in these studies could in part be an artefact, as the quality of the single indicator used in these studies has in that case changed over time. Chapter 5 will test whether or not the relative contributions of the different resources have changed over time.

Chapter 6 will investigate the relationship between inequality during the process through which education is attained (IEOpp) and inequality in the outcome of that process (IEOut). These two types of inequality provide complementary information, but the current literature fails to take this into account. In order to make the best use of this complementarity, one would need to move beyond separately presenting estimates of IEOpp and IEOut and towards an integrated analysis of the two. Chapter 6 will present such an integrated analysis by showing that a method commonly used for estimating IEOpps proposed by Mare (1981) also implies a decomposition of IEOut as a weighted sum of IEOpps, where the weights are a substantively meaningful function of the probabilities of passing the different transitions between levels of education. This decomposition also makes it possible to study the effect of educational expansion on IEOut.

Chapter 7 will present a way to deal with an influential critique by Cameron and Heckman (1998) on the estimates of IEOpp proposed by Mare (1981). Cameron and Heckman (1998) argued that these estimates measure the effect on the average probabilities of passing from one educational level to the next within groups defined by the observed variables rather than the causal effects of these variables on an individual's probability of passing. Moreover, they showed that these group level effects are different from the individual level effects, but that in the literature the group level effects are often interpreted as individual level causal effects. The easiest solution to this discrepancy is to interpret the results of the model proposed by Mare (1981) as group level effects. Alternatively, one could try to estimate individual-level effects. This is, however, much more difficult, as one would also need to control for the heterogeneity between respondents due to unobserved variables (Cameron and Heckman, 1998; Allison, 1999; Mare, 1993). In this chapter I will propose one possible solution, which is to perform a sensitivity analysis by formulating a set of scenarios that vary in the amount of heterogeneity between respondents due to unobserved variables, and estimate the individual-level effects within each of these scenarios. Such a sensitivity analysis will give an idea of the plausible range of individual-level effects.

The final chapter will discuss the extent to which the original research question can be answered and what each of the chapters contribute to what was already known about the trend in the inequality of access to education in the Netherlands. Some of the limitations of the studies collected in this dissertation will also be discussed and some of the areas where this type of analysis can be further improved will be identified.

Chapter 2

Trends in educational inequality in the Netherlands:
A replication and a critique

2.1 Introduction

The degree to which a person's success in education is predetermined by family background is often regarded as the most important indicator of the extent to which a society's resources are distributed based on merits rather than on ascribed statuses. Historical changes in this pattern of achievement versus ascription are therefore of eminent importance. Fortunately, changes over time in educational attainment can be properly monitored by comparing (synthetic) cohorts. Persons born in the same year are likely to enter the schooling system at the same point in time, and the rather rigid nature of formal schooling will ensure that most persons from the same cohort will be subjected to approximately the same educational arrangements. Using cohort comparisons, even a single cross-sectional survey with data on the respondents' education and their family background will contain enough information to enable a historical trend in educational inequality over a period of approximately 40 years to be studied. Many previous studies have enhanced this design by combining data from multiple surveys held at different points in time. Such pooling of cross-sectional surveys leads to larger sample sizes, and thus more statistical power, but also makes it possible to study longer periods of time by combining recent and older surveys covering cohorts that are no longer or not yet available in a single dataset. Also, by continuing to use older surveys, research in this tradition has found a natural way of incorporating past insights into current research, thus facilitating true accumulation of knowledge.

This chapter will continue this tradition by replicating and updating a well-known study on the Netherlands of this kind, conducted by De Graaf and Ganzeboom (1993). These authors combined data from 10 surveys held between 1970 and 1987 covering cohorts born between 1891 and 1960, thus firmly establishing the historical rise of educational mobility (i.e. downward trends in effects of parental status) for the Netherlands. In this replication, I will add data from another 33 surveys. These additional surveys add approximately 60,000 observations, and thus considerable more precision, but also contain information on more recent periods (adding cohorts born between 1960 and 1980), thus making it possible to study the trend for a longer period of time. The surveys used by De Graaf and Ganzeboom (1993) and in this replication are listed in the appendix to this chapter. The analysis will be guided by the following two questions:

> To what extent has there been a historical trend towards less inequality in educational opportunities and in educational outcomes between persons from different status backgrounds?
>
> To what extent do the conclusions by De Graaf and Ganzeboom (1993) hold when using more, and more recent data?

There are two reasons for choosing the study by De Graaf and Ganzeboom (1993) as a benchmark. First, it was part of a much-cited collection of studies of trends in inequality of educational attainment in 13 different countries (Shavit and Blossfeld, 1993) and stood out at the time because of its deviant results: the Netherlands, together with Sweden, was the only country that reported a substantial change towards less inequality of educational attainment. Second, it examined both the association between the highest achieved level of education and family background (Inequality of Educational Outcome, or IEOut) and the association between the probabilities of passing transitions between levels of education (Inequality of Educational Opportunity, or IEOpps), and found a trend towards more mobility in both, while many other studies tend to report only on one of these. IEOpp, which represents inequality during the process of attaining education, and IEOut, which represents inequality in the final outcome of the educational attainment process, are both of substantive interest and complement one another. While subsequent research (e.g. Ganzeboom and Luijkx, 2004b) has already examined the additional available data from the Netherlands in passing, there has not been a major update of the De Graaf and Ganzeboom findings since 1993.

This chapter will not only replicate De Graaf and Ganzeboom (1993) using more data, but it will also critique and improve some of the methods used by these authors. The criticism will come in two parts. First, the 1993 study contains some errors that can be easily rectified within the current context. These errors and their consequences will be discussed during the replication. Second, I will point out that the methods used by De Graaf & Ganzeboom — and replicated in this chapter — do not make the best use of the available information, and I will suggest five improvements. These five improvement require either the estimation of new models, or a substantial reevaluation of the interpretation of the existing models, and each will be discussed in a separate subsequent chapter in this dissertation. The nature of these possible improvements will be further introduced in the conclusions of this chapter.

This chapter will continue with a brief description of the structure of the Dutch educational system, followed by a review of a score of previous empirical studies on trends in inequality of educational attainment in the Netherlands, and in particular a detailed synopsis of De Graaf and Ganzeboom (1993), the benchmark study that will be replicated. Next, the design of the replication will be discussed by introducing the added data, followed by the results of the empirical analysis. This chapter will provide conclusions and the five suggestions for making better use of this type of data, and introduce the subsequent chapters of the dissertation.

2.2 The Dutch education system

The Dutch education system has been subject to a number of developments and reforms. A uniquely important watershed was the introduction of the 'Mammoet Wet' or 'Mammoth Law' in 1968, that established the structure shown in Figure 3.1. This reform is important to most studies in this dissertation because it was implemented at about the middle of the observation period. This means that there are plenty of observations before and after this reform, so any effect it may have had should be clearly visible in these studies. It is convenient to choose this system as a reference and translate all other systems in terms of this reference. The basic structure of the system at that point can be sketched as follows. Primary education (LO) started at about age 6 and took 6 years. After finishing LO, a person must choose between four programmes at the secondary level: LBO (junior vocational education), MAVO (junior general secondary education), HAVO (senior general secondary education), and VWO (pre-university education). Then there are three pathways available if you wish to continue to more advanced levels of education. LBO and MAVO give access to MBO (senior secondary vocational education). HAVO gives access to HBO (higher vocational education). VWO gives access to WO (university). However, students can deviate from these three standard paths, for instance by choosing to 'move up' within their current column (LBO to MAVO, MBO to HBO, and so forth), or 'move down' in the next column (HAVO to MBO, and VWO to HBO).

It is important to note that the Mammoth Law left some features of Dutch education intact. In particular, it did not tinker with the age at which children move on from primary to secondary education. Throughout the period of study, the basic cut-off point in Dutch education has been at age 12, after 6 years of compulsory primary education[1]. This transition — which almost always implied, and still does imply, a transfer to a different school environment — has been a stable feature. By contrast,

[1]Throughout most of this period of study pre-primary education or kindergarten for children aged four and five was also quite common, but not compulsory. It became compulsory for childer aged five in 1985.

Figure 2.1: The Dutch education system after 1968

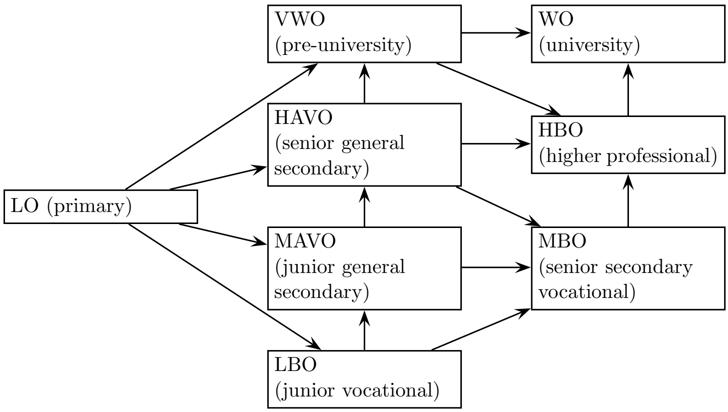

the Mammoth Law changed the existing structure in many other ways, some dramatic, others more cosmetic. One major reform was that the Mammoth Law encouraged schools to offer programmes at different levels (LBO, MAVO, HAVO, VWO) in the same institution and also to offer a common and comprehensive first year (the 'bridge year'), thus giving the opportunity of postponing the decision concerning which secondary level programme to enter by another year. Among the programmes, the HAVO level was new, although it resembled in some respects a programme that had been phased out in 1968 that was exclusively accessible to girls (MMS). The 6-year VWO programme assembled several previously existing older programmes (some lasting 5 years) that gave direct access to university (WO) at age 18. In addition to the comprehensive 'bridge year', moving between programs after the choice had been made was made easier.

A somewhat cosmetic aspect of the Mammoth Law was that it changed the names of most of the programmes. Table 2.1 shows the programmes with their Mammoth names, together with the equivalent old names, the number of years of education they involve, their British-language equivalents, and their ISCED classification (UNESCO, 1997).

Table 2.1: Conversion of old educational levels into new educational levels and simplified educational levels

English name	before 1968	after 1968	duration[a]	ISCED
primary	LO	LO	6	1
extended primary	VGLO	-	7	1
junior vocational	LTS /ambachtschool	LBO	10	2C
junior vocational	LHNO / huishoudschool	LBO	10	2C
junior general secondary	ULO / MULO	MAVO	9 / 10	2B[b]
senior secondary vocational	MTS	MBO	14	3C
senior general secondary	MMS	HAVO	11	3B[b]
pre-university	HBS	VWO	12	3A[b]
pre-university	lyceum	VWO	12	3A
pre-university	gymnasium	VWO	12	3A
higher professional	HTS	HBO	15	5B
university	universiteit	WO	16	5A

[a] Years refer to the situation after 1968 except VGLO.

[b] These levels were originally intended to be terminal levels of education for most students (so 2C or 3C) but evolved into levels that primarily grant access to subsequent levels of education.

2.3 Previous research

A summary of the results of all studies assessing trends in inequality in educational attainment using a (pooled) cross-section design[2] in the Netherlands is shown in Table 2.2. The first to apply the cohort design in the Netherlands for the study of changes in educational inequality were Peschar et al. (1986) and Peschar (1987). These authors used data from a single survey (net82n, see the appendix to this chapter) and found no change over cohorts in the association between the highest achieved level of education and family background, the IEOut. The studies by Peschar and colleagues were followed by Ganzeboom and De Graaf (1989) and De Graaf and Ganzeboom (1990), who improved on the earlier work by assembling multiple surveys. As a consequence these studies contain much more observations and cover a long period of time. These two studies and all subsequent studies using a similar design have found a downward trend in IEOut, suggesting that Peschar's earlier finding of no trend was a matter of lack of statistical power.

A key feature of these early studies is that they examine the association between the highest achieved level of education and family background, in other words, they look at IEOut instead of IEOpp. This can be justified as it is the highest achieved level of education that influences later life chances, so it is inequality in the highest achieved level of education that ultimately influences inequality in other domains of life. However, the focus on final level completed has been criticized by Mare (1981) for not modelling the process through which education is attained. Mare argued that attaining a final educational level consists of a sequence of steps between levels, called transitions, and that the causal effects of parental background exert their influences at those transitions and not directly on the highest achieved level of education. Moreover, Mare (1981) showed that the IEOut is not a mere average of patterns of inequality at separate transitions, but that it is heavily influenced by the distribution of education. This is an important finding because over cohorts the educational distribution changes dramatically, so that any change in the effect of parental background on highest achieved

[2]The main alternative to this pooled cross-sections design is constituted by studies using panel data that follow a cohort of students through their educational career. Examples of this type of study are (Peschar, 1978; De Jong et al., 1982; Bakker et al., 1982; Meesters et al., 1983; Vrooman and Dronkers, 1986; Faasse et al., 1987; Dronkers and Bosma, 1990; Bakker and Schouten, 1991; Dronkers, 1993; Bakker and Cremers, 1994; Rijken et al., 2007). Unlike the studies using cross-sectional surveys the cohort-panel studies find at best mixed evidence for a declining trend in IEOpp at this transition. Panel studies have the advantage that one can study actual transitions between levels of education, and thus get better estimates of IEOpps than is possible using cross-sectional data. However, these data usually cover only the early part of the educational career, making them ill-suited for studying IEOut. In fact, most of these studies focus exclusively on the transition between primary and secondary education when students choose their initial secondary programme instead of the entire educational career. Moreover, they cover a relatively short period of historical time, and within this period the trend is usually estimated by comparing a small number of cohorts (often only two).

level of education could be due to changes in the distribution of education rather than by true causal changes of the inequality in the process of educational attainment. As a consequence, Mare proposed to model the effects of social background on the transition probabilities instead of on the highest achieved level. This model is known by a variety of names, including the sequential response model (Maddala, 1983), the continuation ratio logit (Agresti, 2002), the model for nested dichotomies (Fox, 1997), or simply the 'Mare' model (Shavit and Blossfeld, 1993). This article will use the term 'sequential logit model' (Tutz, 1991) to emphasize that logistic regression is used to model the probabilities of passing transitions.

Those studies in Table 2.2 that use OLS, LISREL, scaled-association models and log-linear models measure IEOut, while studies using the sequential logit model estimate IEOpp. The findings of these studies can be summarized as strong evidence for a linearly declining trend in IEOut and a linearly declining trend in the IEOpp involving the choice of whether or not to continue after primary education, but only weak evidence for a negative trend in IEOpp involving the choice of further enrolment after completing lower levels of secondary education, and no evidence for a trend in IEOpp involving the choice to finish tertiary education.

Table 2.2: Results concerning trends in IEOpp and/or IEOut in the Netherlands from previous studies

study	parental background[a]	birth cohorts	method	trend	linear
Peschar et al. (1986)	fed	1925–1964	log-linear	no trend	
Peschar (1987)	fed	1925–1964	log-linear	no trend	
Ganzeboom and De Graaf (1989)	fed	1891–1960	log-linear	negative	
De Graaf and Ganzeboom (1990)	fed	1891–1960	log-linear	negative	yes
De Graaf and Ganzeboom (1993)	foc fed	1891–1960	OLS & sequential logit	mixed[b]	yes
De Graaf and Luijkx (1995)	foc	1917–1957	OLS	negative	yes
Ganzeboom et al. (1995)	foc fed	1908–1968	OLS	negative	
Ganzeboom (1996)	foc fed	1920–1965	OLS & sequential logit	mixed[c]	yes
Wolbers and de Graaf (1996)	foc fed med	1928–1967	sequential logit	no trend	
Rijken (1999)	foc	1900–1965	OLS & sequential logit	negative	
Korupp et al. (2000)	foc moc	1927–1975	LISREL	mixed[d]	
Sieben et al. (2001)	foc fed med	1925–1974	LISREL	mixed[e]	yes
Korupp et al. (2002)	foc fed moc med	1923–1962	OLS	negative	
Gesthuizen et al. (2005)	foc fed med	1923–1978	survival	negative	

[a] foc = father's occupational status, fed = father's education, moc = mother's occupational status, med = mother's education

[b] Negative for effect on highest achieved level of education and for the first transition, negative but not significant for the second transition, not negative for the third transition

[c] Negative trend for effects on highest achieved level of education and effect of father's occupation on transition from primary education versus more education and effect of father's education on transitions from primary education versus more education and lower secondary versus more education

[d] Significant difference in effect of father's occupation on daughter's education between cohorts 1927–1958 and 1959–1975, all other effects show no trend.

[e] Significant negative trend in effect father's education, no significant trend in father's occupation or mother's education

2.3.1 Results from the study by De Graaf and Ganzeboom (1993) and the design of the replication

De Graaf and Ganzeboom (1993) looked at the changes in effect of father's education and father's occupation on the offspring's highest achieved level of education (IEOut) and the probabilities passing three transitions (IEOpps). The transitions De Graaf and Ganzeboom (1993) analysed were: 1) from no diploma to any diploma in secondary education (LBO, MAVO, HAVO, MBO, and VWO) or higher, 2) from any diploma in lower secondary education (LBO, MAVO) to any diploma in higher secondary (HAVO, MBO, and VWO) or tertiary education, 3) from any diploma in higher secondary education to completed tertiary education (HBO, and WO). The historical trends were assessed by comparing seven ten-year wide cohorts that were born in 1891–1960. To evaluate trends, the authors tested whether differences between cohorts can best be summarized by a single linear trend instead of a separate estimate for each cohort. The findings can be summarized as follows:

1. Inequality of Educational Outcomes

   (a) The data are better described by a linear main effect of cohort and by linear trends in the effect of the father's education and the father's occupational status than by separate estimates for every ten-year wide cohort.

   (b) The effects of father's education and father's occupational status both decrease over time.

   (c) Father's education has a stronger impact than father's occupational status, and the effect of the father's occupational status declines faster than the effect of father's education. As a consequence, the effect of father's education increases relative to father's occupational status.

2. Inequality of Educational Opportunities

   (a) There has been a negative linear trend for both the effects of the father's education and the effect of the father's occupational status on success at the first transition, between primary and secondary education.

   (b) There has also been a negative linear trend for both the effects of father's education and the effect of father's occupational status on the second transition, from lower-level secondary programmes to completing higher-level secondary programmes that give access to programmes at the tertiary level. However, this trend is non-significant, except for the effect of father's education for men.

(c) There is no trend in the effects of the father's education and the father's occupational status on the third transition.

The data to be used in this replication have been taken from 55 surveys held in the Netherlands that were harmonized as part of the International Stratification and Mobility File [ISMF] (Ganzeboom and Treiman, 2009). All ISMF surveys contain information on gender, age (year of birth), the highest achieved level of education and the occupational status of the father (foc). Some of these surveys also contain additional information about mother's occupational status (moc), and father's and mother's highest achieved level of education (fed and med). The appendix to this chapter reports for each survey the year in which it was held, the birth cohorts covered by the survey, the number of respondents, which additional variables are available, and whether or not it was used by De Graaf and Ganzeboom (1993). In order to replicate the analysis by De Graaf and Ganzeboom (1993) only the ISMF surveys that also contain information about the father's education will be used. The number of such surveys available in the ISMF has increased from 10 in the 1993 study to 43 in this replication. The number of respondents has increased from 6,128 men and 5,116 women to 35,846 men and 34,022 women. This replication also covers more recent birth cohorts: 1891–1980 instead of 1891–1960.

In order to replicate the approach followed by De Graaf and Ganzeboom (1993), only respondents who were older than 25 at the time of the interview were used in the analysis, but no upper age limit was imposed. The lower limit ensures that the respondents have finished their full-time education and so their final highest achieved level of education is known. The absence of an upper age limit makes it possible to include the earliest cohort, 1891–1900, whose members were at least 62 when they were interviewed in 1958, when the earliest ISMF survey for the Netherlands was held. A concern might be that including data from older respondents can cause selection on the dependent variable, as higher educated people are more likely to live longer than lower educated people. Such a selection on the dependent variable can lead to biased estimates of the effect of explanatory variables (Breen, 1996). For this reason the earliest cohort is excluded from the analysis in the subsequent chapters. However, in order to match the design of De Graaf and Ganzeboom (1993), this cohort will be included in this chapter.

Education of parents and respondents were measured in four categories: primary education (LO), lower secondary education (LBO and MAVO), higher secondary education (HAVO, MBO, and VWO), and tertiary education (HBO and WO). Notice that the second transition groups together two very different choices: HAVO and VWO are immediately chosen after primary education, while MBO can only be chosen after having finished lower secondary education. Also HAVO and VWO are not intended as terminal levels of education, while MBO is a terminal level of education. However, these levels were grouped together because not all surveys distinguished between them. In concordance with the study by De Graaf and Ganzeboom (1993), the four levels were given the numerical values 1 to 4. Using these quantifications, the distribution of the respondent's highest achieved level of education over cohort and gender is displayed in Figure 2.2. It shows that people who were born more recently are more likely to have completed higher secondary or tertiary education and much less likely to have completed only primary education. This increase in average level of education across cohorts is found in many — if not all — countries, and is usually referred to as 'educational expansion' (Hout and DiPrete, 2006). Figure 2.2 shows that educational expansion in the Netherlands occurred later for women than for men. Both the initial disadvantaged position of women and the decline, or even reversal, of this disadvantage are also features commonly found in other countries (Hout and DiPrete, 2006).

Figure 2.2: Distribution of highest achieved level of education

Father's occupational status was measured according to the father's score on the International Socio-Economic Index of occupational status [ISEI] (Ganzeboom and Treiman, 2003) which was originally measured on a continuous scale from 10 (low status) to 90 (high status), but has been rescaled here to a range between 0 and 8.

2.4 Results

2.4.1 Inequality of Educational Outcome

To model Inequality of Educational Outcome, a linear regression of highest achieved level of education was estimated separately for men and women. The effects of the father's occupation and the father's education capture the IEOut. These effects are allowed to vary over cohorts by adding interactions with either a set of dummy variables for the birth cohorts (to capture a non-linear trend) or a single metric variable (to constrain the trend in IEOut to be linear). This results in a set of nested models, which are presented in panel (a) of Table 2.3 together with their R^2. These models are compared using nested F-tests. These F-tests compare two models, a larger and a smaller model, in the situation that the smaller model can be obtained by imposing a linear constraint on the larger model. The R^2s of the two models being compared and the F-statistic are related to one another according to the following formula:

$$F = \frac{(R^2_u - R^2_c)/df_{num}}{(1 - R^2_u)/df_{denom}} \tag{2.1}$$

R^2_u stands for the R^2 of the larger (unconstrained) model, R^2_c represents the R^2 of the smaller (constrained) model, df_{num} represents the numerator degrees of freedom or the number of linear constraints, and df_{denom} the denominator degrees of freedom or the number of observations minus the number of parameters in the larger model[3].

There are two aims to these comparisons. The first aim is to assess whether trends in the effects of father's occupation and education are linear. This is based on the comparison of the models in which cohort is represented by a set of dummy variables with the models in which cohort is represented by a linear trend. The second aim is to assess whether there has been any trend at all. This conclusion can be made by comparing the models without a trend interaction term with models with a linear trend. A problem with the approach by De Graaf and Ganzeboom (1993) is that they started their analysis by imposing the constraint that the main effect of cohort is linear. Once the main effect of cohort is constrained to be linear, this can influence the linearity of the interaction terms (the trends in the effects of father's education and occupation). This would be unfortunate since it is these latter trends that are of primary interest; they are the trends in IEOut we are testing. It is safer to leave the trend in the intercept free to vary, while testing the trends in the effects. This appears to matter, as the original sequence of tests by De Graaf and Ganzeboom (1993) leads to a linear effect of father's education for women and non-linear trends in all other effects, while in

[3] De Graaf and Ganzeboom (1993) erroneously state that the denominator degrees of freedom equals the number of observations minus the number of parameters in the smaller model.

the sequence preferred here only the effect of father's education for men is non-linear. However, a graphical comparison of the estimates using separate cohorts and a linear trend as in Figure 2.3 shows that in all cases the linear trend provides a reasonable summary of the changes over cohorts.

Table 2.3: Test for trends in Inequality of Educational Outcome

(a) Fit statistics

model	constraints father's education	father's occupation	Intercept	number of parameters	R^2 men	women
sequence of models as used by De Graaf and Ganzeboom (1993)						
1	dummies	dummies	dummies	26	0.277	0.366
2	dummies	dummies	trend	19	0.276	0.365
3	trend	dummies	trend	12	0.276	0.365
4	trend	trend	trend	5	0.275	0.363
5	trend	constant	trend	4	0.274	0.362
6	constant	trend	trend	4	0.270	0.360
7	trend	trend	constant	4	0.240	0.319
preferred sequence of models						
8	trend	dummies	dummies	19	0.277	0.365
9	trend	trend	dummies	12	0.276	0.365
10	trend	constant	dummies	11	0.275	0.365
11	constant	trend	dummies	11	0.272	0.362

(b) Tests

contrast	df_{num}	men df_{denom}	F	p	women df_{denom}	F	p
sequence of models as used by De Graaf and Ganzeboom (1993)							
2 - 1	7	35813	7.066	0.000	33999	5.499	0.000
3 - 2	7	35820	2.763	0.007	34006	0.978	0.445
4 - 3	7	35827	3.262	0.002	34013	13.840	0.000
5 - 4	1	35834	59.733	0.000	34020	40.670	0.000
6 -4	1	35834	260.915	0.000	34020	142.971	0.000
7 - 4	1	35834	1743.814	0.000	34020	2355.229	0.000
preferred sequence of models							
8 - 1	7	35813	1.940	0.059	33999	0.647	0.717
9 - 8	7	35820	4.177	0.000	34006	1.048	0.395
10 - 9	1	35827	58.644	0.000	34013	42.788	0.000
11 - 9	1	35827	186.240	0.000	34013	179.970	0.000

Figure 2.3: Inequality of Educational Outcome (unstandardized coefficients)

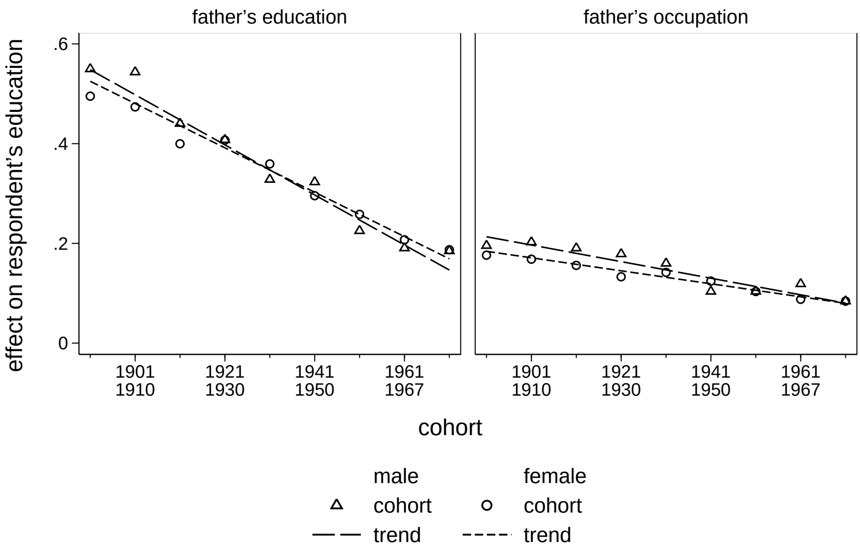

The parameter estimates of models 4 and 9 are presented in Table 2.4. The difference between these models is that in model 4 the trend in the intercept is linear, while in model 9 it is left free to vary across cohorts. The main effects of the father's education and the father's occupation represent the IEOut in the earliest observed cohort, 1891–1900. As in De Graaf and Ganzeboom (1993), these effects are not standardized, so the effect of father's education is the effect of an increase in father's education by one level, while the effect of father's occupation is the effect of an increase in father's occupational status by 1/8th of the range of the occupational status scale. The trend parameters are changes in these effects per decade. One way to get a sense of the size of the trend is to extrapolate when the IEOut will have completely disappeared if the trend continues unchanged. According to model 9, the effect of father's education will have completely disappeared for the cohort that will be born in 2009[4] and 2017 for men and women respectively. Similarly, the effect of father's occupation will have disappeared for the cohort born in 2025 and 2041 for men and women respectively.

De Graaf and Ganzeboom (1993) also claim to have found that, in relative terms, the effect of father's education has become more important than father's occupation.

[4]The effect of father's education for men in model 9 is .547 - .050 $\times$ t, this will be zero at t= -.547/-.050 = 10.94 decades after 1900, that is in 2009

Table 2.4: Estimates of IEOut and trend in IEOut

	men		women	
	model 4	model 9	model 4	model 9
effect father's education	0.578	0.547	0.491	0.525
	(29.94)	(27.05)	(27.45)	(28.26)
trend in effect father's education	-0.056	-0.050	-0.038	-0.045
	(-16.15)	(-13.65)	(-11.96)	(-13.42)
effect father's occupation	0.214	0.213	0.182	0.184
	(18.14)	(18.07)	(16.63)	(16.82)
trend in effect father's occupation	-0.017	-0.017	-0.013	-0.013
	(-7.73)	(-7.66)	(-6.38)	(-6.54)

t statistics in parentheses

The intercept and the dummies for the different cohorts are not reported

They explained this finding by assuming that father's occupation corresponds more closely to the economic resources available in a family while the father's education correspond more closely to the cultural resources in the family. The decrease in the influence of economic resources would be in line with modernization theory, while cultural reproduction theory would predict an enduring influence of the cultural resources of the parents on especially secondary and tertiary education.

However, their analysis of this issue is problematic for two reasons. A first concern arises because they make the effects of father's education and occupational status comparable by standardizing within each cohort, and provide no justification for this choice. This method of standardization implies that the value of the respondent's education changes as the distribution of the respondent's education changes, and that the values (in terms of being able to influence their offspring's education) of the father's education and occupational status change as the distributions of these variables change. The first idea is common, and is often referred to as diploma inflation. However, the parameterization chosen by De Graaf and Ganzeboom overlooks the fact that the value of a level of education is not only determined by how many people have a certain diploma, but also by the demand for people with that diploma. For this reason, the simpler parameterisation of standardizing between cohorts is preferred here, i.e. standardizing using the overall standard deviations of the variables instead of using the cohort-specific standard deviations.

A second, and more serious, concern is that De Graaf and Ganzeboom use the model with linear trends in the effects of father's education and occupation to compute the ratios of these effects. The assumption of linear trends implies changing ratios unless there is no trend in both effects or when both effects are 0 at cohort 0. So this model is clearly not appropriate for studying changes in the relative sizes of two effects. The appropriate model is to estimate separate effects for each cohort without imposing a linear change over time (model 1 in Table 2.3). Figure 2.4 shows how these ratios change over cohorts according to the different models and standardizations. The preferred ratios are those based on coefficients that were standardized between cohorts in model 1, the bottom right graph of Figure 2.4.

Unlike the conclusions of De Graaf and Ganzeboom (1993), the size of the effect of the father's education relative to the father's occupation seems to actually decline, instead of rise. There is however one feature of this trend that is hard to explain, and that is the sudden spike in the ratio for men from the cohort 1941–1950. In other data, such a spike would be attributed to outlying observations, or — as this dataset consists of multiple surveys — an outlying survey. However, this cohort happens to be the largest cohort containing the largest number of observations and surveys, so that no single observation or survey can have a major influence. This feature thus remains unexplained.

Figure 2.4: The effect of the father's education relative to the father's occupation in model 1

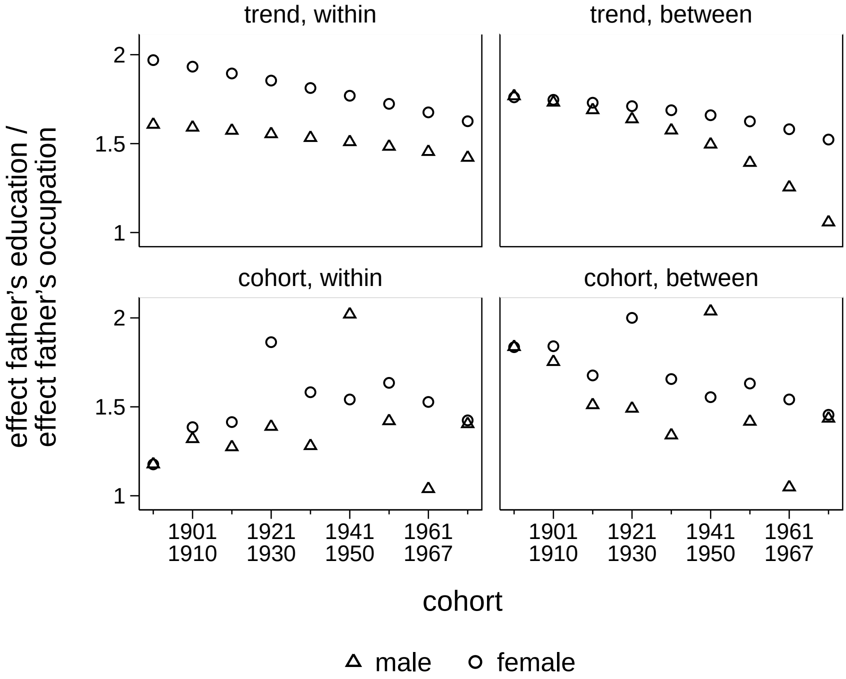

2.4.2 Inequality of Educational Opportunity

As in De Graaf and Ganzeboom (1993), the IEOpps are defined as the association between father's occupational and educational status and the probabilities of passing three transitions: 1) from a diploma in primary education to any diploma in secondary or tertiary education, 2) from a diploma in lower secondary education to any diploma in higher secondary or tertiary education, 3) from any diploma in higher secondary education to completed tertiary education. These IEOpps were measured using the sequential logit model as proposed by Mare (1981). Separate logit models were estimated for each transition, conditional on having passed the previous transition. As with the analysis of IEOut, the analysis of IEOpp will consist of two parts: a sequence of tests on the trends in the effects of the family background variables, and a comparison of the effects of father's education and father's occupation by computing the ratios of standardized coefficients. The concerns with the approach taken by De Graaf and Ganzeboom discussed when analysing IEOut also apply here: (A) it is better not to constrain the trend in intercept to be linear before testing whether the trend in the effects of family background variables is linear; (B) when standardizing, it is better to standardize between cohorts and not within cohorts, and (C) when comparing ratios of effects across cohorts, it is better to base those ratios on a model that allows the effects to change freely across cohorts.

Tables 2.5, and 2.6 represent the tests for the trend in IEOpp equivalent to the tests performed on the trend in IEOut. Instead of comparing the models using the F-test, the models are compared using the likelihood ratio test, as the F-test is only available for models that are estimated using ordinary least squares. The difference between the F-test and the likelihood ratio test is that the F-test takes into account the fact that it is based on a finite sample (through the denominator degrees of freedom) while the likelihood ratio test assumes an infinitively large sample (Long, 1997). Since the sample size is very large, the distinction is negligible in this case. The test statistic of the likelihood ratio test is twice the absolute value of the difference in log likelihood of the two models that are compared, and is χ^2 distributed if the null hypothesis is true.

Despite the enormously expanded database, the results are very similar to the ones found by De Graaf and Ganzeboom (1993), as can be seen in Table 2.7 and Figure 2.5. There is a clearly declining trend in IEOpp for the first transition, and there is still mixed evidence for a trend in IEOpp for the second transition. The trend at the third transition is more complex: the effect of father's education for men is significantly declining, while the effect of father's occupational status for women is significantly increasing. The latter increase in inequality could be due to the fact that the group of women at risk of entering tertiary education has become a lot less selective over the period that is being studied, meaning that there is more room in the recent cohorts for an effect of family background. Also the IEOpps are highest in the first transition, and lowest in the last transition. This pattern has been identified by De Graaf and Ganzeboom (1993) and in many other countries (Hout and DiPrete, 2006), and two explanations have been put forward by Mare (1980). First, the higher transitions are usually made when the person is older, and older persons are less dependent on their family than younger persons. Second, there is only a selected sub-sample at risk of making the higher transitions - those who passed the previous transitions - and this selection causes a negative correlation between unobserved and observed variables, leading to an underestimation of the effects of the observed variables. Using pooled cross-section data from a single country, little can be said about the relative merits of these two explanations (but see Rijken, 1999).

Table 2.5: Tests for trend in Inequality of Educational Opportunity, fit statistics

model	constraints father's education	father's occupation	Intercept	number of parameters	log likelihood Transition 1 men	women	Transition 2 men	women	Transition 3 men	women
sequence of models as used by De Graaf and Ganzeboom (1993)										
1	cohort	cohort	cohort	26	-11459.63	-11466.26	-17627.82	-15974.66	-13113.89	-10400.2
2	cohort	cohort	trend	19	-11466.91	-11474.63	-17637.98	-16005.31	-13131.26	-10402.6
3	trend	cohort	trend	12	-11472.44	-11480.77	-17643.93	-16007.09	-13137.93	-10405.32
4	trend	trend	trend	5	-11481.91	-11498.8	-17654.31	-16125.8	-13160.26	-10428.09
5	trend	constant	trend	4	-11502.24	-11534.05	-17655.57	-16136.99	-13161.88	-10434.22
6	constant	trend	trend	4	-11496.98	-11503.68	-17655.87	-16126.16	-13170.69	-10428.98
7	trend	trend	constant	4	-11789.77	-11974.58	-17705.37	-16268.52	-13160.43	-10428.27
preferred sequence of models										
8	trend	cohort	cohort	19	-11463.52	-11477.02	-17633.94	-15977.43	-13116.76	-10403.12
9	trend	trend	cohort	12	-11473.47	-11478.73	-17641.29	-15980.9	-13124.52	-10408.33
10	trend	constant	cohort	11	-11493.79	-11509.3	-17642.48	-15983.47	-13125.37	-10413.23
11	constant	trend	cohort	11	-11488.98	-11484.27	-17643.89	-15981.32	-13128.64	-10408.98

Table 2.6: Tests for trend in Inequality of Educational Opportunity, Tests

		Transition 1				Transition 2				Transition 3			
		men		women		men		women		men		women	
contrast	df	χ^2	p	χ^2	p	χ^2	p	χ^2	p	χ^2	p	χ^2	p
sequence of models as used by De Graaf and Ganzeboom (1993)													
2 - 1	7	14.543	0.042	16.740	0.019	20.319	0.005	61.285	0.000	34.748	0.000	4.801	0.684
3 - 2	7	11.060	0.136	12.273	0.092	11.892	0.104	3.574	0.827	13.328	0.065	5.435	0.607
4 - 3	7	18.955	0.008	36.066	0.000	20.753	0.004	237.417	0.000	44.674	0.000	45.546	0.000
5 - 4	1	40.659	0.000	70.495	0.000	2.532	0.112	22.371	0.000	3.231	0.072	12.252	0.000
6 -4	1	30.133	0.000	9.754	0.002	3.118	0.077	0.719	0.396	20.861	0.000	1.784	0.182
7 - 4	1	615.704	0.000	951.552	0.000	102.121	0.000	285.432	0.000	0.342	0.559	0.350	0.554
preferred sequence of models													
8 - 1	7	7.779	0.353	21.508	0.003	12.242	0.093	5.527	0.596	5.743	0.570	5.838	0.559
9 - 8	7	19.895	0.006	3.435	0.842	14.690	0.040	6.936	0.436	15.527	0.030	10.423	0.166
10 - 9	1	40.648	0.000	61.130	0.000	2.390	0.122	5.147	0.023	1.691	0.193	9.803	0.002
11 - 9	1	31.024	0.000	11.064	0.001	5.206	0.023	0.857	0.355	8.230	0.004	1.294	0.255

Figure 2.5: Inequality of Educational Opportunity

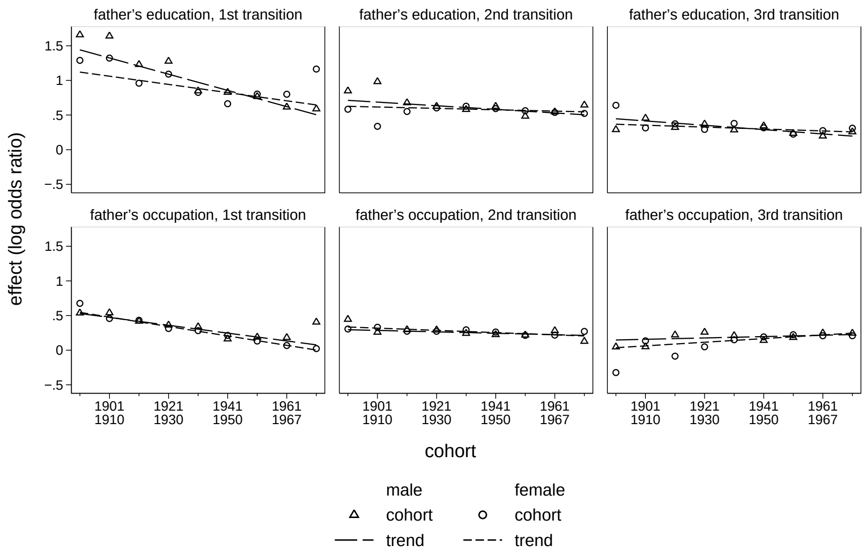

The relative sizes of the effects of father's occupation and father's education can be studied by computing the ratio of the standardized coefficients of these two variables. The results are shown in Figure 2.6. A striking feature of these graphs is the large degree of variability of some of these estimates, so much so that one of these estimates (the youngest cohort for women in the first transition) needed to be truncated in order to obtain interpretable graphs. This degree of uncertainty is understandable: there is very little information present in the data because either there are very few people at risk of passing (transition 3), or virtually everybody passes that transition (transition 1). For this reason there is also little evidence for a trend in the ratio of the effects of father's education and father's occupation.

Table 2.7: Estimates of IEOpp and trend in IEOpp

	men		women	
	model 4	model 9	model 4	model 9
Transition 1				
effect father's education	1.407	1.440	1.104	1.121
	(13.64)	(13.26)	(13.10)	(13.11)
trend in effect father's education	-0.109	-0.117	-0.054	-0.059
	(-5.50)	(-5.50)	(-3.12)	(-3.31)
effect father's occupation	0.530	0.530	0.572	0.545
	(13.23)	(13.20)	(14.56)	(14.04)
trend in effect father's occupation	-0.056	-0.057	-0.073	-0.068
	(-6.39)	(-6.38)	(-8.45)	(-7.85)
Transition 2				
effect father's education	0.676	0.710	0.532	0.628
	(11.50)	(11.59)	(8.91)	(10.33)
trend in effect father's education	-0.019	-0.026	0.009	-0.010
	(-1.77)	(-2.28)	(0.85)	(-0.92)
effect father's occupation	0.297	0.296	0.431	0.335
	(8.34)	(8.21)	(10.84)	(8.45)
trend in effect father's occupation	-0.010	-0.010	-0.033	-0.016
	(-1.59)	(-1.55)	(-4.73)	(-2.27)
Transition 3				
effect father's education	0.536	0.446	0.376	0.367
	(8.80)	(7.20)	(5.29)	(5.08)
trend in effect father's education	-0.049	-0.031	-0.016	-0.014
	(-4.56)	(-2.86)	(-1.33)	(-1.14)
effect father's occupation	0.125	0.148	0.018	0.034
	(2.95)	(3.47)	(0.35)	(0.66)
trend in effect father's occupation	0.013	0.010	0.029	0.027
	(1.80)	(1.30)	(3.50)	(3.13)

z statistics in parentheses

The intercept and the dummies for the different cohorts are not reported

Figure 2.6: The effect of the father's education relative to the father's occupation in model 1

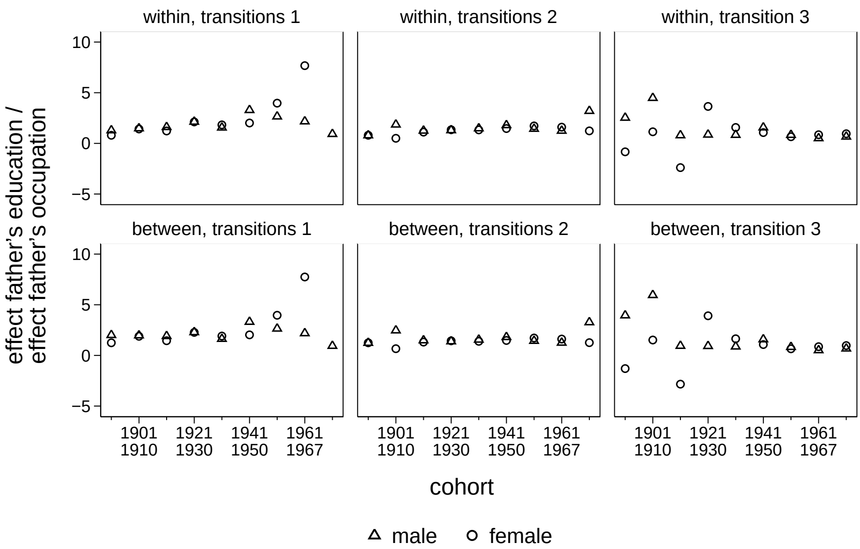

2.5 Summary and discussion

2.5.1 Summary

When studying the effect of parental background on educational attainment, one has to distinguish between two types of effects: the effect on the highest achieved level of education, and the effect on the probabilities of passing the transition between the levels of education that make up the educational system. The former represents the inequality in the end result of the educational process, while the latter represents inequality during the process of attaining education. For this reason they are called Inequality of Educational Outcome (IEOut) and Inequality of Educational Opportunity (IEOpp), respectively. This chapter examined long-term trends in IEOut and IEOpp in the Netherlands by replicating a study by De Graaf and Ganzeboom (1993) using more data and more recent data. This study was chosen as a benchmark as it is much cited and provides estimates of both IEOpp and IEOut. The aim of this replication was to answer the following two questions: (A) To what extent has a trend toward less inequality in educational opportunities and in educational outcomes between persons from different family backgrounds occurred in the Netherlands? (B) To what extent do the conclusions by De Graaf and Ganzeboom (1993) hold when using more and more recent data?

Despite the fact that this replication used a little more than five times as many respondents and covered 20 additional years, the results were largely the same as in the benchmark study: negative trends in IEOut, and in IEOpp for the transition whether or not to continue after primary education, mixed evidence for a negative trend in IEOpp for the choice of track during secondary education, and mixed evidence for trends in IEOpp for the transition whether or not to finish tertiary education. The major deviation from the findings in the benchmark study involved the relative impact of the father's education compared to the father's occupational status. Due to an error in their method, De Graaf and Ganzeboom (1993) concluded that the father's education had become relatively more important, while this replication, using the correct method, found no such trend.

2.5.2 Discussion: how the remaining chapters can improve on this study

The design in this study contain five problems, each of which will be discussed in a subsequent chapter in this dissertation. The first problem is that values need to be assigned to each level of education in order to study IEOut, and the scale of education used in this study is rather crude and arbitrary: 1 for only primary education, 2 for lower secondary education, 3 for higher secondary education, and 4 for tertiary education. More sophisticated *a priori* scales of education exist, mostly based on the institutional number of years assigned to each level. In either case, the value of each level of education is assumed to remain constant over time. This assumption can be questioned as the large increase in the number of people with a higher level of education can be assumed to have led to a decrease in the value of higher levels of education. In Chapter 3 I will empirically estimate a scale of education in order to examine this hypothesis, and to compare the resulting scale with *a priori* scales, including the crude measure used here.

The second problem refers to the way trends in effects are estimated. Two extreme methods were used to test for trends. On the one hand the trend is constrained to be linear, while on the other hand completely separate effects are estimated for each cohort. An intermediate solution is to estimate the trend as a smooth curve. This also allows one to estimate at which point in time such a trend changed. This will be done in Chapter 4 using local polynomial regression. In the other chapters, trends will be estimated using restricted cubic splines, which are more convenient to estimate but less suitable for exactly pinpointing when the trend changed, as the restricted cubic spline model imposes constraints on the change in trend near the beginning and the end of the period under study (Harrell, 2001).

The third problem refers to the informal way in which the hypothesis concerning changes in the relative influence of father's education and father's occupational status were tested. In Chapter 5 I will propose a model that can be used to explicitly test whether the relative contributions of parental education and parental occupational status has remained constant or not. Moreover, this chapter will also investigate whether relative influences of the father and the mother have remained constant or not.

The fourth problem is that the estimates of IEOut and IEOpp are treated separately, while in fact the two are related, since IEOut represents inequality in the end result of the educational process and IEOpp inequality during the educational process. Chapter 6 will explore the way in which these two types of inequality complement one another. IEOut will be shown to be a weighted sum of IEOpps, such that an IEOpp receives more weight if: 1) the proportion of people 'at risk' of making that transition increases; 2) the proportion passing that transition is closer to 50%, that is, passing or failing that transition cannot be described as 'almost universal'; and 3) the difference in expected level of education between those who pass and those who fail to make the transition increases, that is, the expected gain achieved by passing increases. Educational expansion would thus condition the role of IEOpps to predict IEOut by making some transitions become more important, for instance because more people have become at risk of passing that transition, while other transitions have become less important, for instance because virtually everybody passes that transition.

The fifth problem is that the estimates of IEOpp are potentially sensitive to the exclusion of (unobserved) variables, like ability or motivation of the respondent. Excluding these variables from the model will change the results even if these variables are uncorrelated with the variables in the model at the first transition. This means that the estimates of IEOpp are likely to be biased even in the best possible case, when none of the omitted variables are confounding variables. This potential influence is the result of two mechanisms: The first mechanism is that leaving a variable out means that the probabilities will be averaged over these unobserved variables. This will influence the estimates of IEOpp as the IEOpp is a non-linear transformation of these probabilities (it is the logarithm of the ratios of odds). In Chapter 7 I will call this the averaging mechanism. The second mechanism is that the IEOpps at later transitions are based on a selected sample: only those students who are at risk of passing these transitions. This selection can cause a negative correlation between the observed and unobserved variables. I will call this the selection mechanism. Finding a solution to this problem is difficult as such an analysis has two contradictory aims: on the one hand one wants to perform an empirical analysis while on the other hand one wants to control for variation that has not been observed. Chapter 7 proposes one possible solution: estimate the IEOut given a scenario specified by the researcher concerning the unobserved variable. By presenting results of multiple scenarios one can give an indication of the range of plausible values of IEOpp.

Appendix: Description of data sources

Table 2.8: Description of surveys on the Netherlands that are part of the International Stratification and Mobility File (Ganzeboom and Treiman, 2009)

study[a]	year	cohorts	N	additional variables[b]
net58	1958	1891–1933	987	
net67	1967	1896–1942	1162	
net67t	1967	1927–1942	387	
net70[c,d]	1970	1891–1945	1569	fed med
net71c[d]	1971	1898–1944	1223	fed
net71[c,d]	1971	1891–1946	1507	fed
net74p[c,d]	1974	1891–1949	852	fed med
net76j[c,d]	1976	1900–1951	689	fed
net77[c,d]	1977	1891–1952	3116	fed med moc
net77e[c,d]	1977	1891–1952	1339	fed
net79p[c,d]	1979	1891–1954	1344	fed med
net81e[d]	1981	1891–1956	1697	fed
net82e[d]	1982	1891–1957	1184	fed
net82n[c,d]	1982	1917–1957	1783	fed med
net82u[c,d]	1982	1917–1957	621	fed moc
net85o[c,d]	1985	1904–1960	3372	fed med
net86e[d]	1986	1893–1961	1266	fed
net86l[d]	1986	1907–1961	3094	fed med
net87i	1987	1907–1962	1335	
net87j[c,d]	1987	1897–1962	715	fed
net87s	1987	1915–1962	730	moc
net88o[d]	1988	1912–1963	3644	fed
net90[d]	1990	1920–1965	1894	fed
net90o[d]	1990	1913–1965	3471	fed
net91j	1991	1909–1966	736	
net92f[d]	1992	1915–1968	1644	fed med moc
net92o[d]	1992	1911–1967	3690	fed
net92t[d]	1992	1903–1967	1753	fed med moc
net94e[d]	1994	1905–1969	1445	fed med moc
net94h[d]	1994	1913–1969	913	fed med moc

(Continued on next page)

Table 2.8 – continued from previous page

study[a]	year	cohorts	N	additional variables[b]
net94o[d]	1994	1911–1969	3512	fed
net95h[d]	1995	1916–1970	1716	fed med moc
net95s[d]	1995	1925–1970	1688	fed med
net95y[d]	1995	1944–1970	1187	fed med moc
net96[d]	1996	1909–1971	697	fed med moc
net96c[d]	1996	1901–1971	1148	fed med moc
net96o[d]	1996	1911–1971	3823	fed
net96y[d]	1996	1962–1971	271	fed med moc
net98[d]	1998	1902–1973	737	fed med moc
net98e[d]	1998	1908–1973	1314	fed med moc
net98f[d]	1998	1915–1973	1856	fed med moc
net98o[d]	1998	1911–1973	4041	fed med
net99[d]	1999	1906–1974	2150	fed med moc
net99a[d]	1999	1904–1974	7671	fed med moc
net99i[d]	1999	1916–1974	1188	fed med
net00f	2000	1916–1975	1450	
net00s[d]	2000	1930–1975	888	fed med
net02e[d]	2002	1907–1978	1888	fed med moc
net03f	2003	1924–1978	1835	
net03n[d]	2003	1923–1979	7520	fed med moc
net04e[d]	2004	1910–1980	1593	fed med moc
net04i[d]	2005	1912–1980	1540	fed med moc
net06e[d]	2006	1912–1981	1560	fed med moc
net06i[d]	2006	1907–1981	1729	fed med moc

[a] Codes refer to the data references

[b] moc is the mother's occupational status; med is the mother's education; fed is the father's education

[c] used in (De Graaf and Ganzeboom, 1993)

[d] used in replication

Chapter 3

Scaling levels of education

3.1 Introduction

Education is an important stratifying mechanism in modern societies (Hout and DiPrete, 2006). For that reason, education is entered in many models as either an explanans or as the explanandum, often by turning education into a metric variable using institutional durations, in other words, the number of years a 'standard student' would take to obtain a diploma for an educational category. The advantage of this way of scaling education is that it has a meaningful metric and that these values can often be easily obtained from official or pseudo-official documents. However, there are also a number of disadvantages. First, it conflates duration with value, which are two related but different concepts. Second, these scales can sometimes lead to a rank order of educational categories that does not conform to *a priori* knowledge about the educational system, thus requiring *ad hoc* corrections. Finally, this way of scaling education leads to constant values of educational categories over time, while there is an influential hypothesis — the credential inflation hypothesis — that the values of educational categories have changed over time. In order to deal with these limitations, in this chapter I will estimate a new scale of education for the Netherlands in the 20$^{\text{th}}$ century. These levels of education are not directly observed, instead one can observe the respondents' educational category and the association between these categories and a number of positive outcomes, for example: a better job, a higher income, or access to more desirable social networks. This chapter will centre around one such positive outcome: the respondent's occupational status. The idea is to create a metric variable of level of education by assigning values to each educational category such that this metric level of education optimally predicts the respondent's occupational status. Notice that this implies a distinction between the scaling of education, that is, the relative values assigned to each educational category, and the effect of the metric education variable on occupational status.

This scale will be used to answer two questions. The first question is: Which values best represent each educational category in the Netherlands? The estimated values of the educational categories are put into perspective by comparing the estimated values with the values from a commonly used *a priori* scale (Ganzeboom and Treiman, 2009) for the relative distances between educational categories in the Netherlands that is based on institutional durations. The second question is: How have the values of the educational categories changed over time? There are two mechanisms through which the values of the educational categories can change: First, educational systems are often subject to reform. Such reforms may lead to changes in values of the educational categories that are treated as equivalent, either formally or in practice. This means that such an educational category before and after the reform should be treated as two distinct categories. Second, changes in the number of individuals with higher levels of education relative to the demand for highly-educated workers could lead to changes in the values of the educational categories. (Rumberger, 1981; Clogg and Shockey, 1984; Van der Ploeg, 1994; Wolbers, 1998; Hartog, 2000; Groot and Maassen van den Brink, 2000; Wolbers et al., 2001). The credential inflation hypothesis predicts that the supply of highly-educated labor has increased faster than the demand for highly-educated labor, thereby leading to a decrease in the value of all educational categories. However, not all forms of credential inflation (or for that matter its opposite, credential deflation) will influence the scale of education. The reason for this is that the scale of education only measures the relative distances between the educational categories. So, if all educational categories are equally affected by credential inflation, then the relative distances between the categories, and thus the scale, will remain unchanged. Credential inflation will only influence the scale of education if it affects some educational categories more than others.

3.2 Previous research

The two questions will be answered by decomposing the association between the respondents' educational categories and occupational status into a metric scale for the level of education and the effect of the level of education on the occupational status. Changes over time in the association between educational categories and labor market outcomes have already been intensely studied as part of the controversy surrounding credential inflation. Credential inflation is the hypothesis that the number of people with higher levels of education has increased faster than the demand for these people. As a consequence, those with higher levels of education start accepting lower jobs, pushing those who would normally take those jobs further down, thus leading to a decrease in the value of all the categories of education. Most research in this area does not distinguish between the effect of education and the scale of education (Rumberger, 1981; Clogg and Shockey, 1984; Van der Ploeg, 1994; Wolbers, 1998; Hartog, 2000; Groot and Maassen van den Brink, 2000). The most commonly used measure of credential inflation is the incidence of overeducation, defined as having attained a higher level of education than is required for the job. The evidence regarding the changes in the rate of overeducation is rather mixed: on the one hand some studies find an increase in the incidence of overeducation (Rumberger, 1981; Clogg and Shockey, 1984; Wolbers, 1998), while on the other hand a meta-analysis of these studies shows that there is little empirical evidence for such a trend, neither internationally nor in The Netherlands (Groot and Maassen van den Brink, 2000). However, studying the incidence of overeducation provides only a partial answer to the questions that are posed in this chapter, as it conflates the scale of education with the effect of education.

The study that comes closest to distinguishing between the scale of education and the effect is that by Wolbers et al. (2001), who distinguish between what they call structural change, which corresponds to changes in the scale of education, and change in association, which corresponds to changes in the effect of education. But even though Wolbers et al. (2001) make this distinction in theory, in the end they decide not to apply it in their empirical work. Instead of estimating both the scale and the effect of education, and testing whether or not either has changed over time, they *a priori* fixed the values of the educational categories at the percentage of respondents with at least the same level of education. Their argument for not simultaneously estimating a model with a changing scale of education and changing effect of education is that they claim that this model is not identified (Wolbers et al., 2001, p. 12). However, as I will show in section 3.4, this model is equivalent to a model which includes education as a categorical variable and interacts that categorical variable with time, and is thus identified.

3.3 The Dutch educational system

A short description of the Dutch educational system is given in order to put the scale of education that will be estimated in perspective. This discussion of the Dutch educational system will, in part, be framed as a discussion on what happened before and after the introduction of an important educational reform in 1968 called the Mammoetwet or ‘Mammoth Law’. This does not mean that the Mammoth Law is the only educational reform that occurred during the period under study. It merely means that it was the most comprehensive change to the Dutch educational system. The systems before and after 1968 are presented in Figure 3.1. Although this reform represented a significant change in the system, there are also many features that have remained unchanged. The most important of these is that throughout the twentieth century the educational system in the Netherlands remained a so-called tracked system. Immediately after primary education, students have to choose between four tracks: junior vocational (LBO), junior general secondary (MAVO), senior general secondary (HAVO), and pre-university education (VWO). Within the two lower tracks students can choose to continue to senior secondary vocational education (MBO), higher professional education (HBO) is accessible through HAVO and VWO, while university is accessible through VWO. The abbreviations used above are the names of these levels after 1968, which will be used in this chapter as the generic names for these categories unless it is necessary to refer explicitly to the pre-1968 category.

The main differences before and after the reform of 1968 are that it became easier to move between tracks, and that the choice between tracks can be postponed by a year with the introduction of a common and comprehensive first year immediately after finishing primary education, a so-called ‘bridge year’. Regarding the scaling of levels of education, the most important changes are that the Mammoth Law fundamentally changed the nature of at least two levels. First, with respect to lower general secondary education (ULO and MULO prior to 1968 and MAVO after 1968), the Mammoth law formalized and encouraged a practice which had already started: initially (M)ULO was intended to be a terminal level, educating its students for non-manual occupations that require more schooling than primary education. The role of (M)ULO then gradually changed to a level that prepares for MBO. Second, a new level of senior general secondary education was created, the HAVO. A similar senior general secondary program (MMS) did exist prior to 1968, but this was a school for girls and intended to be a terminal level of education. The HAVO is intended to prepare for HBO. Based on these developments one would expect that (M)ULO was more valuable than MAVO, and that MMS had a different value than HAVO though the direction of this difference is less clear.

Figure 3.1: The Dutch education system

(a) Before 1968

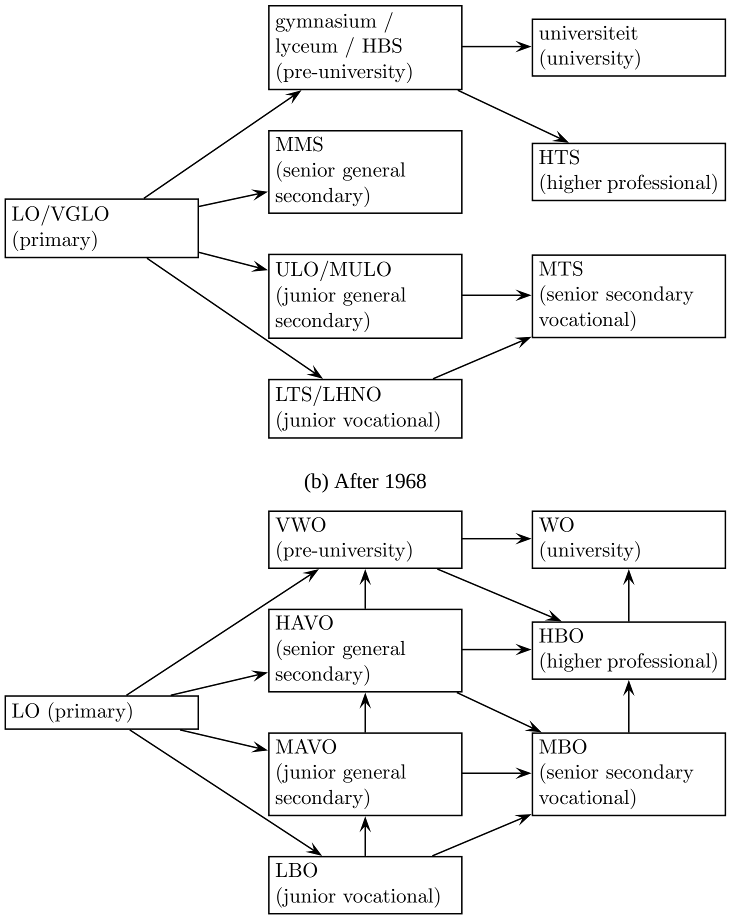

More information about these levels is given in Table 3.1. This table shows the English names for the educational categories, and their Dutch names before and after 1968. In order to get an idea of plausible values of these levels Table 3.1 also reports the institutional duration, the *a priori* scale used in the International Stratification and Mobility File (ISMF) by Ganzeboom and Treiman (2009), and their ISCED classification (UNESCO, 1997). The institutional durations are the number of years a 'normal' student would need to finish this level of education. The *a priori* scale is a measure of the value of each educational category, which uses the institutional duration as a starting point, but applies an *ad hoc* adjustment to make sure that the rank order of each category corresponds to an *a priori* assumption about these values. For the Netherlands this results in an adjustment of the value of MBO. When using institutional years of education, MBO would be assigned a higher value than HAVO and VWO, and is thus ranked above HAVO and VWO. However, obtaining MBO will most likely lead to a blue collar job and obtaining HAVO and VWO will most likely lead to a white collar job, even though both HAVO and VWO are intended as a preparation for further study and not as a preparation for the labor market. For this reason, Ganzeboom and Treiman (2009) apply an *ad hoc* correction by assigning MBO a value between MAVO and HAVO. The metric of the resulting *a priori* scale is called pseudo-years, not only because of this *ad hoc* adjustment, but also because this scale is intended to measure the value of each educational category rather than the duration.

Table 3.1: Conversion of old educational levels into new educational levels

English name	before 1968	after 1968	institutional duration	*a priori* ISMF scale (pseudo-years)	ISCED
primary	LO / VGLO	LO	6 / 7	6	1
junior vocational	LTS / LHNO	LBO	10	9	2C
junior general secondary	ULO / MULO	MAVO	9 / 10	10	2B[a]
senior secondary vocational	MTS	MBO	12 / 14	10.5	3C
senior general secondary	MMS	HAVO	11	11	3B[a]
pre-university	HBS /lyceum / gymnasium	VWO	12	12	3A[a]
higher professional	HTS	HBO	15	15	5B
university	universiteit	WO	16 / 17	17	5A

[a] These programmes were originally intended to be terminal levels of education for most students (so 2C or 3C) but evolved into levels that primarily grant access to subsequent levels of education.

3.4 The model

In this chapter I will scale the educational categories to create a metric education variable in such a way that this metric education variable optimally predicts occupational status. A schematic representation of this model is given in equation (3.1).

$$\begin{matrix}\text{occupational}\\ \text{status}\end{matrix} = \begin{matrix}\text{control}\\ \text{variables}\end{matrix} + \left(\begin{matrix}\text{effect of}\\ \text{education}\end{matrix}\right) \times \left(\begin{matrix}\text{scale of}\\ \text{education}\end{matrix}\right) \tag{3.1}$$

This equation shows that this model will consist of three elements: a set of control variables, optimally scaled education, and an effect of education. A key characteristic of this model is the separation between the effect of the metric education variable and the scaling of the educational categories. In this model it is possible to allow the effect of education to change over one or more variables, for example time, and keep the scaling constant, to keep the effect constant and allow the scaling to change over one or more other variables, allow both the effect and the scaling to change, or keep both the effect and the scaling constant. This model is known under the name: regression with parametrically weighted explanatory variables (Yamaguchi, 2002). It is a special case of the model for estimating a sheaf coefficient (Heise, 1972), which assumes that the effect of the latent variable — in this case scaled education — remains constant. It is also a special case of the Multiple Indicators and Multiple Causes (MIMIC) model (Hauser and Goldberger, 1971) where the latent variable is assumed to be measured with error. Finally, it is also a linear model imposing a proportionality constraint, where the effects of all educational categories are constrained to change by the same proportion.

The simplest version of this model assumes that both the scaling and the effects remain constant, which is equivalent to the model for estimating a sheaf coefficient (Heise, 1972). In this case, the model is just a reparameterization of a model that includes education as a set of dummy variables. The model will be introduced using a simplified example in which there are no control variables present, and only three levels of education are distinguished: primary, secondary, and tertiary, which can be represented as a set of three dummy variables: prim for primary education, sec for secondary education, and tter for tertiary education. Extensions will be added after this basic model has been discussed. The starting point is a linear model estimating the effect of education on occupational status (occ), representing education as a series of dummy variables. Such a model is shown in equation (3.2), wherein the βs are the regression coefficients and ε is a normally distributed error term. In this model, primary education is the reference category.

$$\mathsf{occ} = \beta_0 + \beta_1 \mathsf{sec} + \beta_2 \mathsf{ter} + \varepsilon \tag{3.2}$$

An unconventional way to interpret model (3.2), but not a new way, is that it simultaneously estimates the scale of a single metric variable representing the level of education, and the effect of this metric variable. A scale of educational levels will measure the relative distances between the educational categories. Such relative distances need two constraints: one to fix the origin of the scale and another to fix the unit of the scale. So, if the value of primary education is fixed to 0 and that of tertiary education to 1, then this will fix the origin at primary education and this will fix the unit at the distance between primary and tertiary education. The scaling will assign the position of secondary education relative to these two levels. This new variable (ed) can be written like equation (3.3):

$$\mathsf{ed} = \underbrace{\gamma_1}_{0} \mathsf{prim} + \gamma_2 \mathsf{sec} + \underbrace{\gamma_3}_{1} \mathsf{ter} \tag{3.3}$$

Whereby, the γs define the scale. The effect of education on occupation can be written as in equation (3.4), whereby the effect of this scaled education is called λ_1.

$$\begin{aligned} \mathsf{occ} &= \beta_0 + \lambda_1 \mathsf{ed} + \varepsilon \\ &= \beta_0 + \lambda_1 (\underbrace{\gamma_1}_{0} \mathsf{prim} + \gamma_2 \mathsf{sec} + \underbrace{\gamma_3}_{1} \mathsf{ter}) + \varepsilon \\ &= \beta_0 + \lambda_1 \gamma_2 \mathsf{sec} + \lambda_1 \mathsf{ter} + \varepsilon \end{aligned} \tag{3.4}$$

All parameters in model (3.4) can be calculated from the parameters in model (3.2):

$$\begin{aligned} \lambda_1 &= \beta_2 \\ \gamma_1 &= 0 \\ \gamma_2 &= \frac{\beta_1}{\beta_2} \\ \gamma_3 &= 1 \end{aligned}$$

Model (3.4) is thus just a reparameterization of model (3.2), and does not add anything to the model other than an alternative interpretation of the results. This implies that there is no way to test whether a model that separates the effect from the scale is to be preferred over a model consisting only of a set of dummies, as these two models are equivalent. However, this changes when one allows the effect of education to change over other variables while constraining the scaling to remain constant. This implies a testable constraint. This is illustrated by extending the simplified example to allow the effect of education to change over the variable year. The test of an hypothesis involves the comparison of two models, a constrained model and an unconstrained one. The constrained model is represented in equation (3.5), while the unconstrained model includes interaction terms of year with all the dummies as in equation (3.6).

$$\text{occ} = \beta_0 + (\lambda_1 + \lambda_2 \text{year})(\underbrace{\gamma_1}_{0} \text{prim} + \gamma_2 \text{sec} + \underbrace{\gamma_3}_{1} \text{ter}) + \beta_1 \text{year} + \varepsilon \tag{3.5}$$

$$\begin{aligned} \text{occ} &= \alpha_0 + \alpha_1 \text{year} + \\ &\quad \alpha_2 \text{sec} + \alpha_3 \, \text{year} \times \text{sec} + \\ &\quad \alpha_4 \text{ter} + \alpha_5 \, \text{year} \times \text{ter} + \varepsilon \end{aligned} \tag{3.6}$$

To facilitate the comparison of the two models, equation (3.5) can be rewritten as equation (3.7):

$$\begin{aligned} \text{occ} &= \beta_0 + \beta_1 \text{year} + \\ &\quad \lambda_1 \, \gamma_2 \text{sec} + \lambda_2 \, \gamma_2 \, \text{year} \times \text{sec} + \\ &\quad \lambda_1 \text{ter} + \lambda_2 \, \text{year} \times \text{ter} + \varepsilon \end{aligned} \tag{3.7}$$

If the constrained model is true, then $\alpha_2 = \lambda_1 \gamma_2$, $\alpha_3 = \lambda_2 \gamma_2$, etc. This implies that

$$\begin{aligned} \frac{\alpha_2}{\alpha_4} &= \frac{\lambda_1 \gamma_2}{\lambda_1} = \gamma_2 \\ \frac{\alpha_3}{\alpha_5} &= \frac{\lambda_2 \gamma_2}{\lambda_2} = \gamma_2 \end{aligned} \tag{3.8}$$

In other words, the constraint that needs to be imposed on equation (3.6) in order to get equation (3.7) is $\frac{\alpha_2}{\alpha_4} = \frac{\alpha_3}{\alpha_5}$, and it is this constraint that is being tested. This is a proportionality constraint: the effects of educational categories are allowed to change over time, but the proportional distance between the effects are forced to remain equal. The most convenient way of testing this constraint is by comparing the constrained model and the unconstrained model using a likelihood ratio test. Both models are estimated by assuming that ε is normally distributed with a mean of zero and a constant variance. This is a linear regression in the case of the unconstrained model. The constrained model needs to be estimated using maximum likelihood.

This model can be further extended in several ways: first, the model can easily accommodate more than three levels of education, by adding more level dummies. Second, the effect of scaled education can change over more than one variable. Third, the values assigned to each educational category, that is, the scaling of education, can be allowed to change over one or more variables. For instance, one can allow an educational category to have different values before and after an educational reform, and test whether these values are different. Fourth, one can include control variables. This model and all these extensions are implemented in Stata (StataCorp, 2007) as the `propcnsreg` package (Buis, 2007a), which is documented in Technical Materials I.

3.5 The data

The model requires data on the respondent's occupational status, the respondent's educational category, and three additional sets of explanatory variables: the control variables, the variables along which the scaling of educational categories is allowed to change, and the variables along which the effect of education is allowed to change. These variables are:

- control variables
  - gender of the respondent,
  - potential experience (age minus institutional duration of education),
  - year in which the survey was held,
  - father's occupational status,
  - two-way interactions of father's occupational status and gender, father's occupational status and year of survey, potential experience and gender, and potential experience and year of survey.
- variables along which the scaling of educational categories is allowed to change
  - whether or not a respondent belongs to the pre-Mammoth or the post-Mammoth cohort, defined as the cohort that was 12 years old before and after 1968 respectively,
- variables along which the effect of education is allowed to change
  - year in which the survey was held, and
  - the gender of the respondent.

The data used in this chapter consists of 54 Dutch surveys that were harmonized as part of the International Stratification and Mobility File (Ganzeboom and Treiman, 2009). The surveys are listed in the appendix to this chapter and described in the data references. Only respondents older than 27 and younger than 65 where used in the analysis. This dataset contains 72,666 respondents who meet this criterion and have complete information on all the covariates. Figure 3.2 shows how these observations are distributed across time. It is important to note that information on the early years is based on only a few points in time.

Figure 3.2: Number of observations per year

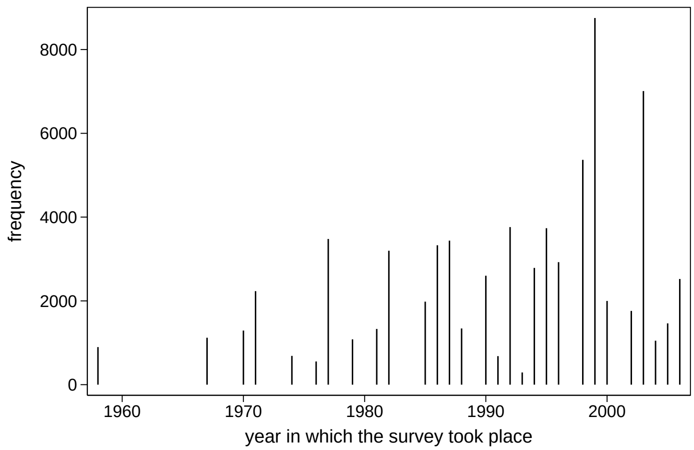

The dependent variable is the occupational status of the most recently held occupation, thus it includes homemakers, unemployed, and early retirees who have had a job in the past. The occupations were scaled to represent occupational status according to the International Socio-Economic Index of occupational status [ISEI] (Ganzeboom and Treiman, 2003), which was originally measured on a continuous scale from 10 (low status) to 90 (high status), but is rescaled here to a range between 0 and 1.

The educational category is measured as the highest category attained by the respondent. The eight categories are defined as in Table 3.1 and will be referred to by their post-1968 names. However, some surveys merged some of the educational categories into one or more ‘combined categories’. Table 3.2 shows how common this practice has been: a majority of surveys have at least one combined category. The most commonly combined category is HAVO/VWO. This is partly due to the fact that MMS is treated here as the pre-1968 equivalent of HAVO, but it was such a small category that earlier surveys routinely merged that category with pre-university edu-

Table 3.2: Prevalence of combined and not-combined educational categories in the data

educational category	number of surveys	number of respondents
not-combined		
LO	54	13,414
LBO	48	16,773
MAVO	46	9,908
MBO	46	14,763
HAVO	24	1,747
VWO	31	1,550
HBO	50	13,668
WO	51	6,962
combined		
HAVO/VWO	27	4,498
LBO/MAVO	5	1,476
HBO/WO	3	395
HAVO/VWO/MBO	2	478
VWO/MBO	1	511
MAVO/MBO	1	199
LBO/MBO	1	144
MAVO/HAVO	1	88

cation (VWO). An attractive characteristic of the method used here for estimating the scale of education is that it can accommodate surveys with combined educational categories without having to combine the categories from the other surveys, thus using the maximum amount of detail available from each survey. This is done by simply treating these ‘combined levels’ as a separate level whose value needs to be estimated, which can be done by adding dummy variables for the ‘combined levels’. A more parsimonious way of dealing with these ‘combined levels’ is by constraining their value to be equal to the average value of their constituent levels. This constraint will also be tested.

The control variables used while predicting the respondent’s occupational status with the respondent’s education are: father’s occupational status, the respondent’s gender, the respondent’s (potential) years of labor force experience, and the year in which the survey was held. Father’s occupational status is measured — just like the respondent’s occupational status — in ISEI scores that have here been rescaled to range between 0 and 1. The year in which the survey was held is included as an approximation of the period in which the respondents held their occupation. This variable ranges from 1958 to 2006. However, as shown in Figure 3.2, the information for the earlier years is rather sparse. The potential experience in the labor market is approximated using age minus institutional years of education. Time and experience are allowed to have non-linear effects by entering them in the model as restricted cubic splines (Harrell, 2001). This means that the range of time and experience is split up at locations called knots. Experience was given knots at 10, 25 and 35 years of potential experience, and year was given knots at 1980, 1990, and 2000. In the sections after the first knot and before the last knot, third-degree polynomials are estimated. These curves are forced to meet at the knots and have the same first and second derivative at that point. The curve is restricted to be linear before the first knot and after the last knot. This model has the advantage of leading to a smooth curve that is more stable than an (unrestricted) cubic splines (Harrell, 2001). The restricted cubic spline, as used in this chapter, is implemented in Stata 10 (StataCorp, 2007) in the `mkspline` command.

The effect of education is allowed to change over time and gender. Time is represented by the same restricted cubic spline as was used for the control variables. The values of the educational categories are allowed to change depending on whether a respondent belongs to the 'pre-Mammoth' cohort or the 'post-Mammoth' cohort. These cohorts are defined as whether or not the respondent was 12 years old before or after 1968. This is a rather crude measure as some respondents were already in a 'Mammoth-like' system before 1968 because the law was preceded by a large number of experiments. However, the data do not contain a more precise measure of which respondent was educated in which system.

3.6 Results

Eight models are estimated and are described together with their fit statistics in Table 3.3. These models differ from one another in the following ways. Models labeled (a) assume that the values of the educational categories remained constant apart from possible changes introduced by the educational reform in 1968, which corresponds to imposing the proportionality constraint. The models labeled (b) allow the values to change over time and between men and women, which corresponds to entering education as a categorical variable and adding interaction terms of each educational category dummy variable with time and gender. Models 1, 2, and 3 differ from one another with respect to which educational categories changed in value in 1968. Model 1 assumes that all categories changed in value, model 2 assumed that only MAVO and HBO changed in value, while model 3 assumed that none of the values changed in 1968. Model 4 forces the value of the combined categories to be equal to the average

Table 3.3: Fit statistics

model	proportionality constraint	scale of category changes in 1968	value of combined category	df	log-likelihood	BIC
1(a)	yes	all	freely estimated	44	29804.36	-59104.53
1(b)	no	all	freely estimated	101	29947.26	-58762.87
2(a)	yes	MAVO, HBO	freely estimated	38	29803.71	-59170.47
2(b)	no	MAVO, HBO	freely estimated	77	29911.04	-58959.34
3(a)	yes	none	freely estimated	36	29775.72	-59136.87
3(b)	no	none	freely estimated	69	29873.37	-58973.63
4(a)	yes	MAVO, HBO	average	30	29767.00	-59186.66
4(b)	no	MAVO, HBO	average	54	29834.32	-59063.60

Table 3.4: Test proportionality constraint

contrast[a]	BIC difference
1(a):1(b)	341.66
2(a):2(b)	211.12
3(a):3(b)	163.24
4(a):4(b)	123.06

[a] The model numbers refer to Table 3.3

Table 3.5: Model selection

contrast[a]	hypothesis	BIC difference
1(a):2(a)	no change in value of LO, LBO, HAVO, VWO, and WO	65.93
2(a):3(a)	no change in value of all categories	-33.59
2(a):4(a)	values of combined categories constrained to mean	16.20

[a] The model numbers refer to Table 3.3

value of their constituent categories, while models 1, 2, and 3 freely estimate those values.

The resulting eight models are compared in Tables 3.4 and 3.5. Table 3.4 gives for each model the test of the proportionality constraint, that is, whether the scale of education has remained constant over time, and Table 3.5 compares the four models with a proportionality constraint against one another. Table 3.4 shows that the differences in the Bayesian Information Criterium (BIC) score[1] is much more than 10 points in favor of the constrained model, which provides "very strong" (Raftery, 1995) or "decisive" (Jeffreys, 1961) evidence in favor of the proportionality constraint. An advantage of BIC differences over tests like the likelihood ratio test is that tests will pick up ever smaller deviations from the null hypothesis as the sample size increases. This is consistent with the logic behind statistical testing, but it also means that statistical tests will pick up substantively irrelevant deviations from the null hypothesis when the sample becomes very large. The comparison of BIC scores avoids this problem. Given that the sample size in this case is approximately 75,000 respondents, the comparison of BIC scores is preferred.

The first two comparisons in Table 3.5 investigate whether the scaling of education was influenced by the implementation of the Mammoth Law in 1968. The first row shows that no evidence was found that the values of LO, LBO, HAVO, VWO, and WO changed before and after the Mammoth law. The second row indicates that there is evidence that the value of MAVO and HBO changed. The third row tests the hypothesis that the combined educational categories can be represented by the average of the values of the constituent categories, instead of estimating a separate value for each combined category. This row shows that the BIC difference supports constraining the values of the combined levels. The preferred model is thus model 4a.

Model 4a separates the effect of education on the occupational status of the respondent from the scale of education. The effects are shown in Figure 3.3, while the scale is shown in Figure 3.4. The effects can be transformed into standardized effects by multiplying them by 1.62, as the standard deviation of the latent education variable is .310 and the standard deviation of the respondent's occupational status is .191. The standardized effects thus range between approximately .5 for women around 1960 and approximately .6 for men around 2005. These are thus sizeable effects. Figure 3.3 also shows that women gain less occupational status from education than men, while the relative values of the educational categories are the same.

The scale of education is presented in Figure 3.4. The bottom two lines show the scale of education as estimated in model 4a, while the top line shows the *a priori* scale. Comparing the estimated scale with the *a priori* scale from the ISMF shows that

[1]The BIC score is computed as: BIC = -2*ln(likelihood) + ln(N)*k, where N is the sample size and k is the number of degrees of freedom.

Figure 3.3: Effect and the trend in the effect of education on occupational status

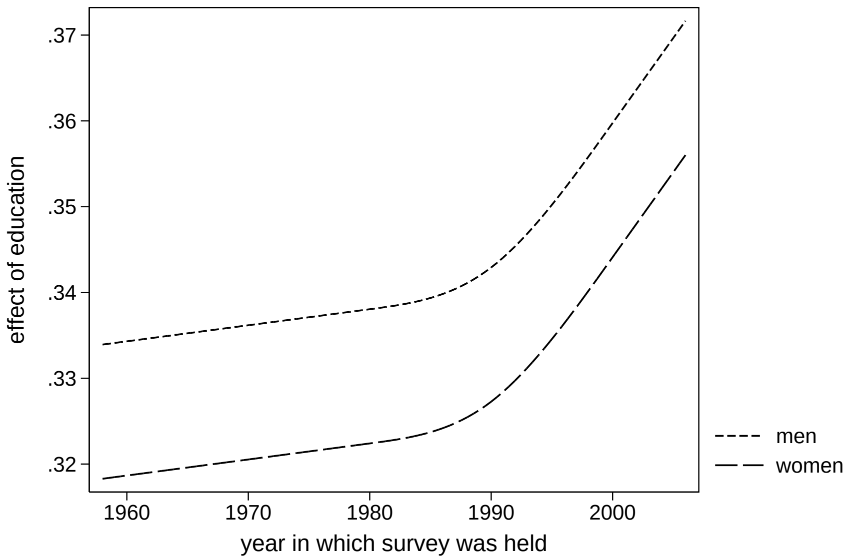

the most striking differences between the two is the value of LBO: LBO is much less valuable than the *a priori* scale suggests. The fact that LBO is the lowest level of secondary education may well result in an extra penalty, explaining why a pseudo-year in LBO is worth less than a pseudo-year in the other forms of secondary education. The value of HAVO and VWO are underrated when using the *a priori* scale of education. This may be explained by fact that some of the respondents with HAVO and VWO as their highest achieved level of education may have started HBO or WO, but never completed it.

Figure 3.4 also shows the comparison between the estimated scale before and after the introduction of the Mammoth Law in 1968. It shows that the value of MAVO decreased after the introduction of the Mammoth Law. Moreover, the rank order changed from a situation where MULO was more valuable than MBO to a situation where MBO was more valuable than MAVO. This is consistent with the transformation from (M)ULO, which was a terminal level in its own right, to MAVO, which is a level preparing for MBO. The decline in the value of HBO may be explained by the fact that the kind of people having access to HBO changed after 1968, as it became accessible through the HAVO.

The numerical values of the *a priori* scale and the estimated scale are presented in

Figure 3.4: Scale of education

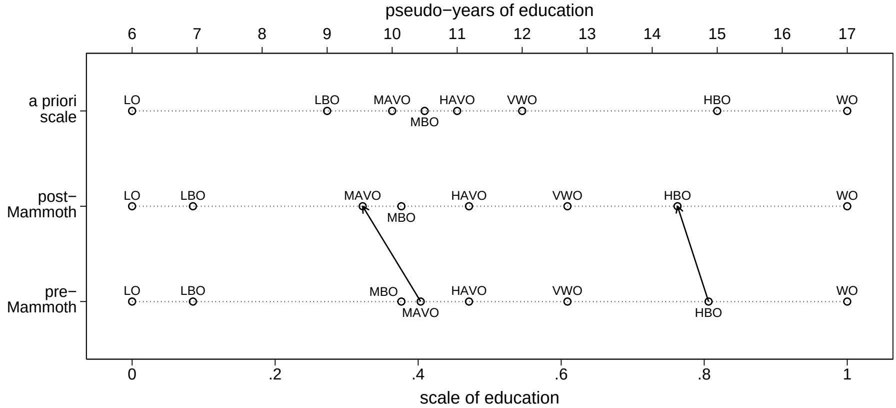

the first two columns of Table 3.6. In the third column, the estimated scale is rescaled such that the metric resembles pseudo-years of education (LO is fixed at 6 and WO is fixed at 17). In the final column, this scale has been stylized by rounding to the nearest half-year. This stylized scale will result in a variable with a metric that is as easy to interpret as the *a priori* scale, but more closely represents education as a resource for attaining occupational status.

Table 3.6: The *a priori* and the estimated scale of education

level	*a priori* scale	estimated scale	rescaled	stylized scale
LO	6	0	6.00	6.0
LBO	9	.085	6.94	7.0
MAVO[a]	10	.404	10.44	10.5
MAVO[b]	10	.324	9.55	9.5
MBO	10.5	.377	10.14	10.0
HAVO	11	.471	11.18	11.0
VWO	12	.609	12.70	12.5
HBO[a]	15	.806	14.87	15.0
HBO[b]	15	.763	14.39	14.5
WO	17	1	17.00	17.0

[a] Before the Mammoth Law: ULO and MULO for MAVO and HTS for HBO

[b] After the Mammoth Law

3.7 Conclusion

This chapter started with the questions concerning which values best represent each level of education in the Netherlands, and how these values have changed over time. Two mechanisms are proposed through which the scale of education could change over time. The first mechanism is educational reform, which can mean that an educational category before and after a reform should be treated as two different categories. In this chapter the focus is on one particular educational reform: the Mammoth Law implemented in 1968. The second mechanism concerns the changes in the supply of highly schooled labor relative to the demand for highly schooled labor. If supply increased (decreased) faster than the demand, then the value of the educational categories is likely to decrease (increase). However, this will only influence the scale of education if the change in value of some categories is stronger than the change in value of other categories, since the scale of education measures only the relative distances between the categories.

In order to study these two issues, a scale of education is estimated such that it is optimal for predicting occupational status, using a model proposed by Yamaguchi (2002), and implemented in the statistical package Stata (StataCorp, 2007) as the `propcnsreg` module (Buis, 2007a) that is documented in Technical Materials I. This model estimates both the effect of education and the scale of education. The model resulted in a scale of education that is summarized in Figure 3.4. This estimated scale was compared with an often-used *a priori* scale as found in the International Stratification and Mobility File (Ganzeboom and Treiman, 2009). The major deviation from the *a priori* scale is that the *a priori* scale overrates the value of LBO, which means that respondents with LBO had on average lower status occupations than was predicted using the *a priori* scale. In order to facilitate the use of this scale in other analyses a stylized version of this scale using the metric of pseudo-years of education was presented in Table 3.6.

Using this model, it was not possible to reject the hypothesis that the introduction of the Mammoth Law in 1968 has not influenced the value of the educational categories for all but two educational categories: MAVO and HBO. The change in the value of MAVO was expected as this level changed from a level that prepared for the labor market to a level that prepared for a subsequent level of education (MBO). A possible reason for the change in the value in HBO could be due to the fact that it became accessible via HAVO. The hypothesis that changes in the supply and demand for highly-educated labor has not led to changes in the relative values of the educational categories could not be rejected. So, the relative distances between the categories remained mostly constant, even though the effect of education on occupational status increased over time.

One way in which the scale could be improved is to use additional indicators like a higher income, and access to more desirable social networks, or one could scale education by how much individuals or families have invested in order to attain a level of education. This would lead to a number of different scales of education. These different scales could be used to create a more comprehensive scale of education by constraining them to be equal. Moreover, by testing whether these scales can be combined into a single scale, one can test the hypothesis that the value of education is a one-dimensional concept rather than a multi-dimensional one. Moreover it may be useful to estimate a scale with higher level of detail, in particular distinguishing between completed and attended educational categories. In the current context this may be most useful for estimating the values of higher general secondary education (HAVO) and pre-university education (VWO). It is likely that a large proportion of respondents that report these categories as their highest achieved level of education have also had some years of higher professional education (HBO) or university, but did not finish these categories. This would lead to an overestimation of the value of HAVO and VWO, as the benefit these respondents received from attending university or HBO is incorrectly assigned to the HAVO or VWO categories. Furthermore, distinguishing between various degrees of incomplete primary education could prove useful when one wants to create a scale that can be used in countries where — or in historical periods when — incomplete primary education is prevalent.

Appendix: Description of data sources

Table 3.7: Merged educational categories and the sizes of the the pre- and post-Mammoth cohorts in Dutch surveys that were post-harmonized in the International Stratification and Mobility File (Ganzeboom and Treiman, 2009)

survey number	survey code[a]	year	cohorts	N pre-Mammoth	N post-Mammoth	merged categories
1	net58	1958	1891–1933	902	0	
2	net67	1967	1896–1942	1,144	0	(LBO MAVO) (HAVO VWO MBO) (HBO WO)
3	net67t	1967	1927–1942			(HAVO VWO) (HBO WO)
4	net70	1970	1891–1945	1,334	0	(LBO MAVO) (HAVO VWO)
5	net71c	1971	1898–1944	1,130	0	(LBO MBO) (HAVO VWO) (HBO WO)
6	net71	1971	1891–1946	1,282	0	(LBO MAVO) (HAVO VWO)
7	net74p	1974	1891–1949	730	0	(HAVO VWO)
8	net76j	1976	1900–1951	669	0	(HAVO VWO)
9	net77	1977	1891–1952	2,659	0	(MAVO MBO) (HAVO VWO)
10	net77e	1977	1891–1952	1,195	0	(HAVO VWO)
11	net79p	1979	1891–1954	1,119	0	
12	net81e	1981	1891–1956	1,518	0	(HAVO VWO)
13	net82e	1982	1891–1957	1,041	0	(HAVO VWO)
14	net82n	1982	1917–1957	1,931	0	
15	net82u	1982	1917–1957	637	0	
16	net85o	1985	1904–1960	3,080	260	(HAVO VWO)
17	net86e	1986	1893–1961	994	110	(HAVO VWO)
18	net86l	1986	1907–1961	2,327	313	(MAVO HAVO) (VWO MBO) (HBO WO)
19	net87i	1987	1907–1962	961	156	(HAVO VWO)
20	net87j	1987	1897–1962	530	72	(HAVO VWO)
21	net87s	1987	1915–1962	620	96	(HAVO VWO)
22	net88o	1988	1912–1963	3,073	654	(HAVO VWO)
23	net90	1990	1920–1965	1,345	425	(HAVO VWO)

Continued on next page

Table 3.7 – continued from previous page

survey number	survey code[a]	year	cohorts	N pre-Mammoth	N post-Mammoth	merged categories
24	net90o	1990	1913–1965	2,704	834	(HAVO VWO)
25	net91j	1991	1909–1966	1,055	404	(LBO MAVO) (HAVO VWO)
26	net92f	1992	1915–1968	1,169	457	
27	net92o	1992	1911–1967	2,793	992	(HAVO VWO)
28	net92t	1992	1903–1967	1,749	790	(HAVO VWO)
29	net94e	1994	1905–1969	776	480	(LBO MAVO) (HAVO VWO MBO)
30	net94h	1994	1913–1969	479	389	
31	net94o	1994	1911–1969	2,687	1,194	(HAVO VWO)
32	net95h	1995	1916–1970	1,096	724	
33	net95s	1995	1925–1970	1,002	595	
34	net95y	1995	1944–1970	39	983	
35	net96	1996	1909–1971	340	247	
36	net96c	1996	1901–1971	794	576	
37	net96o	1996	1911–1971	2,456	1,474	(HAVO VWO)
38	net96y	1996	1962–1971	0	288	
39	net98	1998	1902–1973	364	323	
40	net98e	1998	1908–1973	822	672	
41	net98f	1998	1915–1973	865	891	
42	net98o	1998	1911–1973	2,364	1,885	(HAVO VWO)
43	net99	1999	1906–1974	956	926	
44	net99a	1999	1904–1974	4,274	3,712	
45	net99i	1999	1916–1974	562	570	
46	net00f	2000	1916–1975	678	648	
47	net00s	2000	1930–1975	442	393	
48	net02e	2002	1907–1978	805	901	
49	net03f	2003	1924–1978	741	1,089	
50	net03n	2003	1923–1979	2,753	3,579	
51	net04e	2004	1910–1980	619	728	
52	net04i	2005	1912–1980	612	766	
53	net06e	2006	1912–1981	501	808	
54	net06i	2006	1907–1981	572	891	

[a] Codes refer to the data references

Chapter 4

Deceleration of the trend in inequality of educational outcome in the Netherlands

4.1 Introduction

The association between family socioeconomic status and offspring's educational attainment has been studied long and intensely in social stratification and social mobility research as it is considered a major indicator of the openness of a society (for example Breen and Jonsson, 2005; Hout and DiPrete, 2006). In this chapter I will focus on one aspect of this research theme: the association between socioeconomic background and the highest achieved level of education, which will be called Inequality of Educational Outcome (IEOut), and in particular on how this IEOut has changed over time. The main motivation for studying how IEOut has changed over time is the following dilemma: previous research has found that for the Netherlands IEOut has decreased linearly over time (De Graaf and Ganzeboom, 1990, 1993; De Graaf and Luijkx, 1995; Ganzeboom, 1996; Sieben et al., 2001). Such a linear decrease in the association between socioeconomic background and highest achieved level of education is improbable. A linear trend would eventually lead to a negative association between socioeconomic background and highest achieved level of education, which would mean that having a high status background would become a hindrance instead of an asset for attaining education. This is implausible, and as a consequence the negative trend in IEOut will have to slow down. This leads to the main question that this chapter tries to answer: has there been a deceleration in the trend in IEOut, and if so, when did this deceleration take place? To answer this question, the effect of family background on highest achieved level of education (IEOut) is allowed to change over cohorts according to a smooth but flexible curve. This smooth curve is used to estimate the trend, which is the slope or first derivative of the curve, and the change in trend, which is the slope of the slope or second derivative of the curve. To assess whether and when the negative trend significantly decelerated, I will test whether the change in trend is significantly positive, since a slowing down of a negative trend means that the trend moves from more negative to less negative.

The secondary aim of this chapter is to assess the susceptibility of data assembled in the International Stratification and Mobility File (ISMF) (Ganzeboom and Treiman,

2009) to three potential sources of error. The first potential source of error is due to the fact that the ISMF consists of multiple surveys. The variables of these surveys have been post-harmonized and then stacked in order to create a single dataset. This could lead to a false trend if the quality of the surveys changed systematically over time. Such a systemic change in quality could for example occur because the response rate changed systematically over time. It is likely that the quality of a survey will influence the strength of the associations between the variables in that survey, as the associations in low-quality surveys will be contaminated by more noise than the associations in high-quality surveys. So, a false increasing (decreasing) trend in the association between family background and educational attainment can be expected if the quality of the surveys systematically increased (decreased) over time. The second potential source of error concerns the scale of education as used in the ISMF. This is an *a priori* scale based on the institutional years of education with an *ad hoc* correction of the level assigned to senior secondary vocational education (MBO). In Chapter 3 I proposed a scale with a stronger empirical foundation, which scales the levels such that education optimally predicts occupational status. The former scale will be referred to as the *a priori* scale while the latter will be referred to as the empirical scale. The most prominent difference between these two scales is that the *a priori* scale assigns too much value to lower vocational education (LBO). This can influence the estimated trend in IEOut as this difference in scaling means that the estimated trend is likely to respond differently to changes in the proportion of respondents with lower vocational education over time. The third potential source of error is missing data. This will lead to biased estimates if the likelihood of not answering a question is related to the value of the dependent variable (Allison, 2002). The dependent value in this chapter is the highest achieved level of education, and it is likely that the willingness and ability to finish a survey is associated with the respondent's highest achieved level of education. So, it is plausible that missing data could cause bias in the estimates of IEOut. The severity of this problem is influenced among other things by the proportion of observations that contain missing values. If the proportion of missing values changes over time, then the severity of this problem would change over time and thus also bias the estimated trend. The presence of these three potential sources of error is easier to detect when studying changes in the trend in IEOut, as this is a very subtle analysis. If the potential sources of error matter, then they will certainly show up in such an analysis.

4.2 Previous research

The challenge of studying the trend in IEOut is to cover a sufficient period of time such that the trend, and changes in the trend become visible. A common strategy is to take multiple surveys and compare respondents that are born in different years, that is, the time is captured by comparing so-called synthetic cohorts. By comparing synthetic cohorts, a single survey can cover a period of 40 years (when using respondents who are between 25 and 65 years old). This period can be further extended by adding surveys collected at different times. These cohorts are used as a measure of when the effect of social background on educational attainment took place. This is reasonable, given the strongly age-stratified nature of full-time education, which means that people born in the same year experience a very similar educational system. Within the Netherlands, this technique was first used for the study of the trend in IEOut by Peschar et al. (1986), and has been used in numerous other studies since (Peschar 1987; Ganzeboom and De Graaf 1989; De Graaf and Ganzeboom 1990; De Graaf and Luijkx 1992; De Graaf and Ganzeboom 1993; De Graaf and Luijkx 1995; Ganzeboom et al. 1995; Ganzeboom 1996; Rijken 1999; Korupp et al. 2000, 2002; Breen et al. 2009; and Chapter 2 of the current dissertation), and resulted in the International Stratification and Mobility File (Ganzeboom and Treiman, 2009). Six of these studies (De Graaf and Ganzeboom 1990; De Graaf and Luijkx 1992; De Graaf and Ganzeboom 1993; De Graaf and Luijkx 1995; Ganzeboom 1996; and Chapter 2 of the current dissertation) test whether the trend in IEOut in the Netherlands is linear or not, and none of these studies can reject the hypothesis that IEOut is linearly decreasing over time. In all cases, the tests for non-linearity of the trend were performed by comparing a model with a linear trend with a model with a non-linear trend. This non-linear trend was either a quadratic trend or a discrete trend, where the period was broken up into a series of cohorts, and separate IEOuts were estimated for each cohort. The use of these methods could explain why no non-linearity was found, as quadratic functions can easily be too rigid to adequately represent a non-linear trend, while a discrete trend is very flexible but expends a lot of statistical power, making it hard to find any significant evidence for non-linearity in the trend. The main aim of this chapter is to find out if any non-linearity in the trend can be found if one uses a model that is more flexible than a quadratic function but less flexible, and thus more powerful, than a discrete trend.

Figure 4.1: Number of observations per cohort

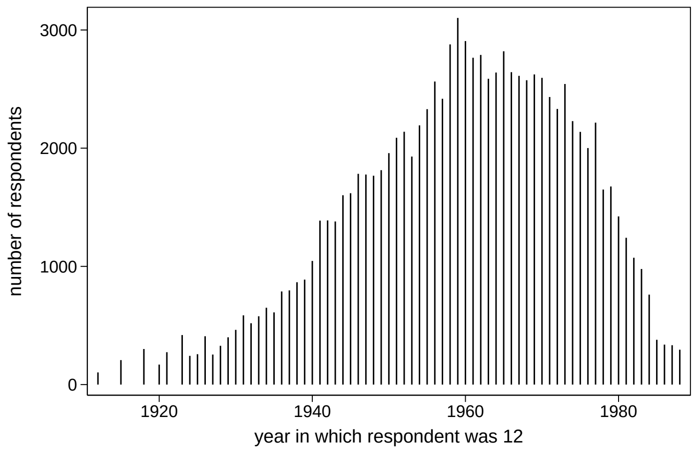

4.3 Data

The data consist of 54 surveys held in the Netherlands between 1958 and 2006 that were post-harmonized as part of the International Stratification and Mobility File (ISMF) (Ganzeboom and Treiman, 2009). Where available, survey weights have been used. The weighted number of respondents is 86,581. The number of respondents is unequally distributed over the cohorts, as is shown in Figure 4.1. Of these respondents, 9,416 lack information on father's occupational status, 651 on the respondent's highest achieved level of education, and 342 on both. If the proportion of missing information varies across cohorts, then this could bias the estimate of the trend. Figure 4.2 shows that the proportion of observations with missing information has changed systematically over time. The reason for these changes across cohorts could be in part an age-effect, as the older cohorts will consist mainly of people that were old at the time of the interview, and in part be a period effect, which can for example capture changes in a general attitude towards surveys and the introduction of computer-assisted interviewing, which makes it harder to skip questions.

Time is measured by the year in which the respondent was 12, which is the age at which most persons in the Netherlands make the most important decision in their educational career. Information is available for the cohorts aged 12 in 1912–1988. IEOut is measured by the strength of the metric association between highest achieved level of education of the respondents and their father's occupational status. Father's occupational status is measured in ISEI scores (Ganzeboom and Treiman, 2003). The original

Figure 4.2: Proportion of observations with missing values per cohort

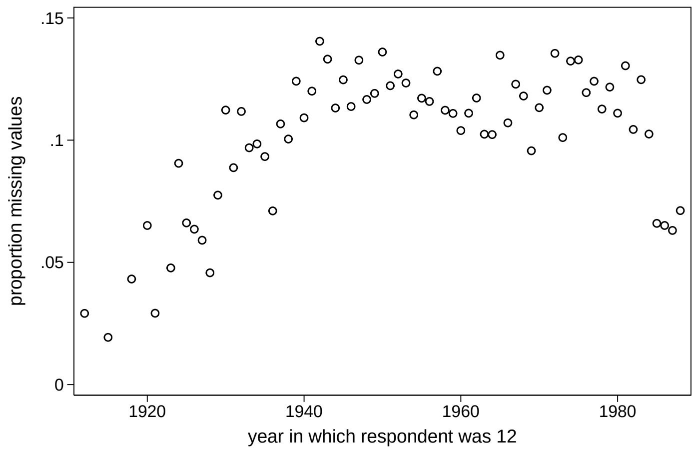

ISEI score is a continuous variable ranging from 10 to 90, but it has been standardized to have an overall mean of 0 and a standard deviation of 1. The highest achieved level of education of the respondents is measured in either the original *a priori* scale from the ISMF or in the empirical scale estimated in Chapter 3. A description of the different levels of education and the two scalings have been reproduced in Table 4.1. The first three columns show the name of each level, their English translation, and their classification in the ISCED (UNESCO, 1997) scheme. The fourth column presents the institutional duration, or the number of years it would take a 'standard student' to finish that level of education. The final two columns present the two scales. The most striking difference between the two scalings is the value of lower vocational education (LBO), whose value is valued considerably higher in the *a priori* scale. Moreover, Table 4.1 shows that in the empirical scale the values of two educational categories, MAVO and HBO, changed in 1968. A major educational reform, the Mammoetwet or 'Mammoth Law', was implemented in that year. The metric of the scales in Table 4.1 is pseudo-years of education, but in the analysis both scales will be standardized to have a mean of zero and a standard deviation of one.

Table 4.1: Conversion of old educational levels into new educational levels and simplified educational levels

level	English translation	ISCED[a]	institutional duration	*a priori* scale	empirical scale
LO	primary	1	6	6.0	6.0
LBO	junior vocational	2C	10	9.0	7.0
MAVO	junior general secondary	2B[b]	9 / 10	10.0	10.5[c] / 9.5[d]
MBO	senior secondary vocational	3C	14	10.5	10.0
HAVO	senior general secondary	3B[b]	11	11.0	11.0
VWO	pre-university	3A	12	12.0	13.0
HBO	higher professional	5B	15	15.0	15.0[c] / 14.5[d]
WO	university	5A	17 / 16	17.0	17.0

[a] (UNESCO, 1997)

[b] These levels were originally intended to be terminal levels of education for most students (so 2C or 3C) but evolved into levels that primarily grant access to subsequent levels of education.

[c] before 1968

[d] after 1968

An important characteristic of this data is that it consists of different surveys. A major advantage of this approach is that this greatly increases the period that can be studied because these surveys were held at different times. This is illustrated in Figure 4.3, which indicates for each survey the cohorts to which it contributes observations. It shows that the oldest cohorts collect their observations from only four surveys, while other cohorts collect their observations from almost all surveys. So peculiarities of individual surveys are most likely to influence estimates in the earliest cohorts, since in these cohorts each survey is responsible for a sufficient proportion of the observations to have a noticeable influence. The characteristics of individual surveys are less likely to have an effect on the estimates in the middle cohorts, as no single survey is dominant in these cohorts. The appendix to this chapter shows that there are considerable variations among surveys in terms of response rate, proportion of missing cases, and the degree of detail in which the variables are measured.

Figure 4.3: Cohorts covered by different surveys (survey numbers correspond to the appendix and are ordered by the year in which the survey was held)

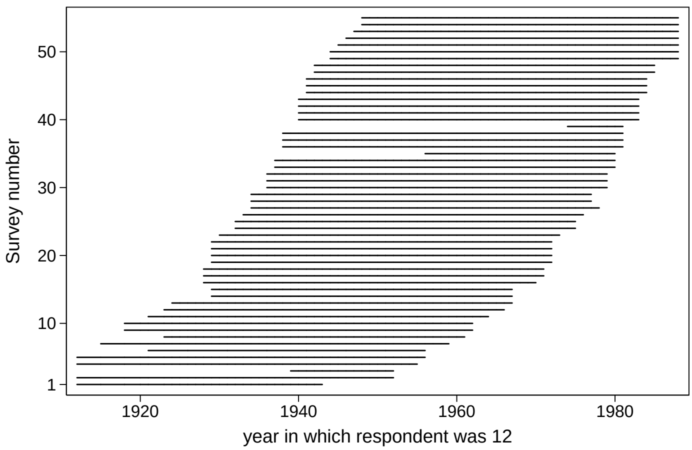

4.4 Method

For the estimation of the non-linear trend, a two-step process has been used. First, a new dataset is created containing, for each annual cohort, an estimate of IEOut for men and women, and their standard errors. The estimates are obtained by regressing the respondent's highest obtained level of education on father's occupational status, separately for men and women and each cohort. An annual cohort is combined with a neighboring cohort if it does not contain enough observations for a stable estimate. This resulted in the following combined cohorts: 1900/1901, 1902/1904, 1905/1907, 1910/1911. Second, a local polynomial curve (Cleveland, 1979; Loader, 1999; Fox, 2000) is fitted through these annual estimates of IEOut. This is done in such a way that estimates with small standard errors, that is, measured with great precision, receive more weight than estimates with large standard errors. These curves also provide estimates of the trend and the change in trend. These are the first and second derivatives of the smooth curve.

An attractive feature of the local polynomial smooth is that it uses information from nearby cohorts to create an improved estimate of IEOut at a cohort. This is illustrated in Figure 4.4. The point on the local polynomial curve for cohort 1938 is computed using the following four steps: First, the observations that will be used in the estimation are selected. This is typically done by selecting a fixed proportion, say 65%, of nearest observations. This is shown in Figure 4.4 in panel (a). Second, observations that are selected are weighted according to their distance from 1938. A common function used to create these weights is the tricube function[1]. The tricube function is shown in panel (b) in Figure 4.4. The tricube weights ensure that cohorts close to 1938 receive more weight when estimating the value of cohort 1938. Third, these weights are adjusted in such a way that they take into account that some cohorts are estimated with much more precision than others. The raw estimates of IEOut are regression parameters, so an estimate of the precision of the estimate is available in the form of the standard error. Weights based on the inverse of the square of the standard error would properly correct for the difference in precision between cohorts. The weights based on proximity to cohort 1938 and the weights based on the precision of the estimates of IEOut are combined by computing the product of these two. This is shown in panel (c) of Figure 4.4. Fourth, a regression of IEOut on cohort, cohort

[1] If the cohort of interest, the 'focal value' is represented by x_0 and the span (half the range that contains 65% of the observations) by h then the value x is assigned the weight

$$W = \begin{cases} \left[1 - \left(\frac{|x-x_0|}{h}\right)^3\right]^3 & \text{if } \frac{|x-x_0|}{h} < 1 \\ 0 & \text{if } \frac{|x-x_0|}{h} \geq 1 \end{cases}$$

squared, and cohort cubed, is estimated using these combined weights. The predicted IEOut from this regression for cohort 1938 is the local polynomial estimate. It uses most information from cohorts that are close by and are estimated with high precision, and less information from cohorts that are far away or are estimated with low precision. Furthermore, the slope of this regression line in 1938 is a local estimate of the trend, and the change in slope in cohort 1938 is a local estimate of the change in trend in 1938. These are obtained by evaluating the first and the second derivatives of the regression line at 1938. The entire local polynomial curve is obtained by repeating this process for all annual cohorts. This procedure is implemented in the `locfit` package (Loader, 2005) in the R statistical computing environment (R Development Core Team, 2005). This also provides confidence envelopes for the curve, the first and the second derivatives, using procedures discussed by Loader (1999).

What makes this method attractive is that it takes an intermediate position between two commonly used alternative methods of estimating a non-linear trend: a quadratic trend and a discrete trend. The first strategy is usually not flexible enough to adequately fit the data. The second strategy is too flexible, which means that too much statistical power is lost, making it hard to find any evidence for a non-linear trend.

The secondary aim of this chapter is to investigate sensitivity of the results to the three potential sources of error: The first potential source of error is the fact that the ISMF consists of multiple surveys that vary in quality. The survey effects are controlled for by adding dummies for surveys, and interacting these dummies with father's occupational status. The dummies are constructed in such a way that the main effect of father's occupational status represents the IEOut in an average survey, so differences between cohorts are no longer the result of differences across surveys. By adding survey dummies, and interactions between the survey dummies and father's occupational status, each survey has its own baseline IEOut, but the trend is constrained to be the same for all surveys. The reason for this choice was that there is good reason to expect that the quality of a survey can influence the effect of father's occupational status on the respondent's education, as the effect is likely to be smaller in more noisy data, but the effect of data quality on the estimated trend of the effect of father's occupational status is much less clear.

Figure 4.4: Obtaining local polynomial regression estimate for IEOut for cohort 1938, adapted from Figure 4.1 in (Fox, 2000, p. 24–25)

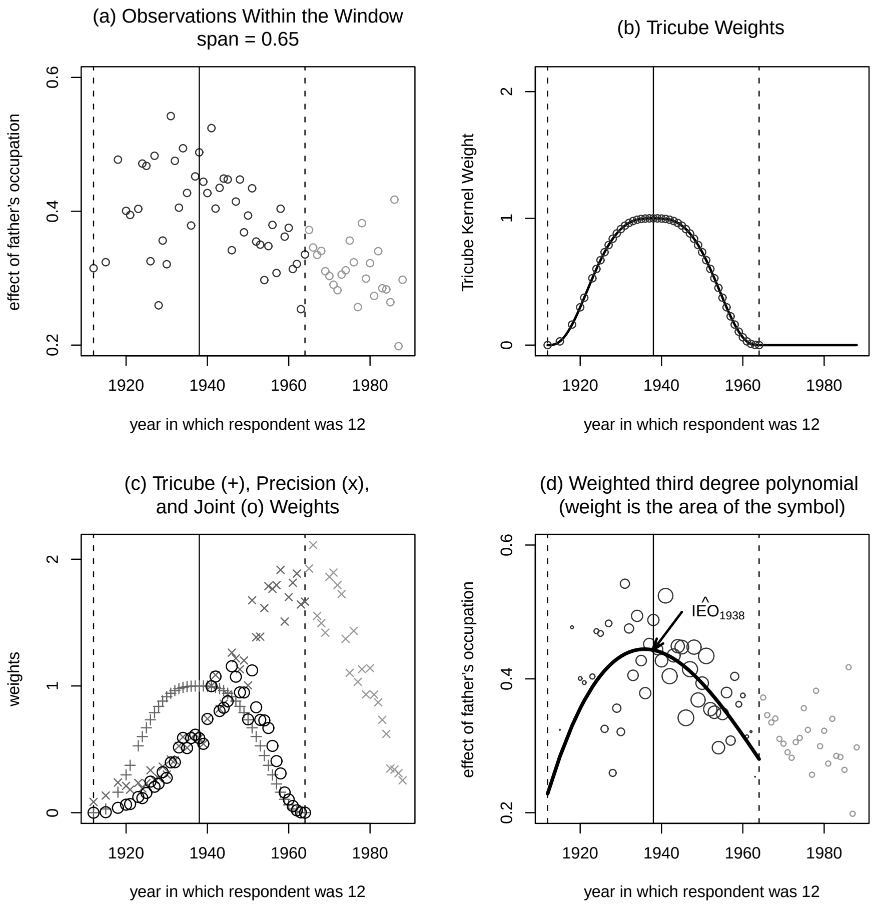

The second potential source of error is the fact that there are multiple ways in which the dependent variable — the respondent's education — can be scaled. The ISMF uses a common strategy by starting with the institutional years of education, the number of years a 'standard student' would need to finish that level of education, and applies an *ad hoc* correction to make sure that the ordering of levels corresponds with an *a priori* ordering. In Chapter 3 I proposed an alternative scale based on the idea that education predicts the occupational status of the respondent, and if education is better scaled then education should be better at predicting occupational status. This way one can estimate an optimal scale of education. These two scales were presented in Table 4.1. By comparing the estimated trend using the *a priori* scale with the estimated trend using the empirical scale, one can assess whether the differences between the scales actually matter.

The third potential source of error is the presence of missing data. The annual estimates of IEOut are controlled for missing data using Multiple Imputation (Little and Rubin, 2002). The idea behind Multiple Imputation is to create multiple 'complete' datasets by first estimating for each missing value a distribution of plausible values, and then drawing multiple values (in this case 5) from this distribution. This is done in Stata (StataCorp, 2007) using the `ice` (Royston, 2004, 2005a,b, 2007, 2009) module. The model of interest is estimated on each 'complete' dataset. The point estimate is the average of the point estimates from the different 'complete' datasets, and the variance of the sampling distribution (the standard error squared) is computed according to equations (4.1) to (4.3) (Little and Rubin, 2002).

$$se^2 = \overline{se^2} + (1 + 1/m)B \tag{4.1}$$

$$\overline{se^2} = \frac{\sum_{j=1}^m se_j^2}{m} \tag{4.2}$$

$$B = \frac{\sum_{j=1}^m (\beta_j - \overline{\beta})^2}{m-1} \tag{4.3}$$

Equation (4.1) shows that the variance of the sampling distribution (se^2) in the case of m 'complete' datasets consists of two elements: $\overline{se^2}$, and B. The first element is described in equation (4.2), and is the average of the variances of the sampling distributions in the different 'complete' datasets. This represents an estimate of the degree of uncertainty about a parameter which ignores the fact that some of the data is itself uncertain as it consists of imputations rather than real observations. The second element, in equation (4.1), and equation (4.3), corrects for this by using the differences in the parameters (β_j) between the different complete datasets as a measure of the uncertainty due to the imputations.

The key issue with multiple imputation is the model used for estimating the imputed values. This model must be at least as flexible as the model of substantive interest (Little and Rubin, 2002). For this reason separate imputation models are estimated for each combination of cohort and survey. Within each of these combinations, imputations for the missing values are created from a model using father's and respondent's occupational status and education, with interactions between whether the respondent is male or female and all these variables. The occupational status of the respondent and the level of education of the father are also used in the imputation model even though they will not be used in the final model of interest, in order to improve the imputations. However, the father's highest achieved level of education is only added when available, which was not the case in 10 surveys. Imputations were only carried out if the cohort-survey combination had at least 20 fully observed cases. As a result, not all missing values were imputed. There were 9,758 missing values for father's occupational status, of which 1,934 could not be imputed, and there were 993 missing values for the respondent's highest achieved level of education, of which 181 could not be imputed. Respondents with missing values that could not be imputed will still be ignored in the analysis.

4.5 Results

The results using estimates of IEOut while controlling for all the potential sources of error and using the empirical scale are shown in detail in Figure 4.5 for men and Figure 4.6 for women. Panels (a) show local polynomial curves fitted through the annual estimates of IEOut with their 95% confidence envelope. The confidence envelopes always remain above zero, indicating that the offspring of fathers with a higher status occupation did, on average, attain more education than the offspring of fathers with a lower status occupation. Panels (a) also show that the level of inequality appears to have changed over time for both men and women. This is tested in the panels (b), which show the trend in IEOut, that is, the first derivatives of the local polynomial curves in panels (a). A period of significant negative trend is found for both men (1941–1960) and women (1952–1977). The hypothesis that the trend is zero in the last period (after 1960 for men and 1977 for women) cannot be rejected, suggesting that the trend has indeed slowed down. Notice however that the confidence envelopes are very wide for both the youngest and the oldest cohorts, so the finding of zero trend in the most recent cohorts could also be due to lack of statistical power. The way to find out if the trend truly decelerated is to also estimate the changes in trends, the second derivatives, which are shown in panels (c). If the trend truly decelerated, then the second derivative should be significantly positive, indicating that the negative trend became less negative. Panels (c) show a significantly accelerating trend (negative second derivative) between 1935 and 1944 for men and 1949 and 1952 for women, but no significant deceleration. For men, the point estimates of the change in trend are positive before the trend became insignificant, providing some indication that the trend decelerated. For women, the point estimate of the change in trend is also briefly positive prior to the trend becoming non-significant, but this period is much shorter, and quickly becomes negative again, so the case for a decelerating trend for women is much less convincing. These results are summarized in panels (d). The curve represents the smooth estimates of IEOut from panel (a), while the shaded areas below that curve represent the periods of significant trend, and the shaded areas above the curve represents periods of significant change in trend.

Figures 4.7 and 4.8 show how controls for the three different potential sources of error influenced these results. Panels (a) and (b) use the *a priori* scale and the empirical scale of education respectively. Panels (c) show the trend using the empirical scale while controlling for survey effects. Panels (d) show the trend using the empirical scale while controlling for missing data. Comparing panels (a) and (b) shows that the scale of education does influence the estimated trend. A decelerating trend was found for men using the *a priori* scale, but this change in trend became insignificant when the empirical scale was used. For women, using the empirical scale leads to a significant positive estimate of the trend prior to the negative trend, and a significant transition between the positive and the negative trend. Neither of these characteristics was present when the *a priori* scale was used. The aspect of the trend that remains largely unaffected by the scale of education is the downward trend during the third quarter of the twentieth century. The panels (c) show the trend in the 'average survey', thus controlling for survey effects. This correction mainly affects the oldest cohorts, since these cohorts contain data from only a few surveys, as was shown in Figure 4.3. As a consequence, a problem in an individual survey could have an influence on the uncorrected results. The younger cohorts contain data from many surveys, so any problems with individual surveys is likely to be averaged out. One important way in which surveys differ from one another is the number of missing values, as is shown in the appendix to this chapter. If this is the main source of differences in the results between models that control and do not control for survey effects then the trends in panels (d), which control for missing data but not for survey effects, should closely correspond to the trends in panels (c). However, the trends in panels (d) closely correspond to the trends in panels (b), indicating that controlling for missing data hardly influences the results.

Figure 4.5: Trend in Inequality of Educational Outcomes and change in trend for men. (IEOut is measured as standardized regression coefficients. The local polynomial smooth has a span of .65 and uses weights proportional to the inverse of the variances.)

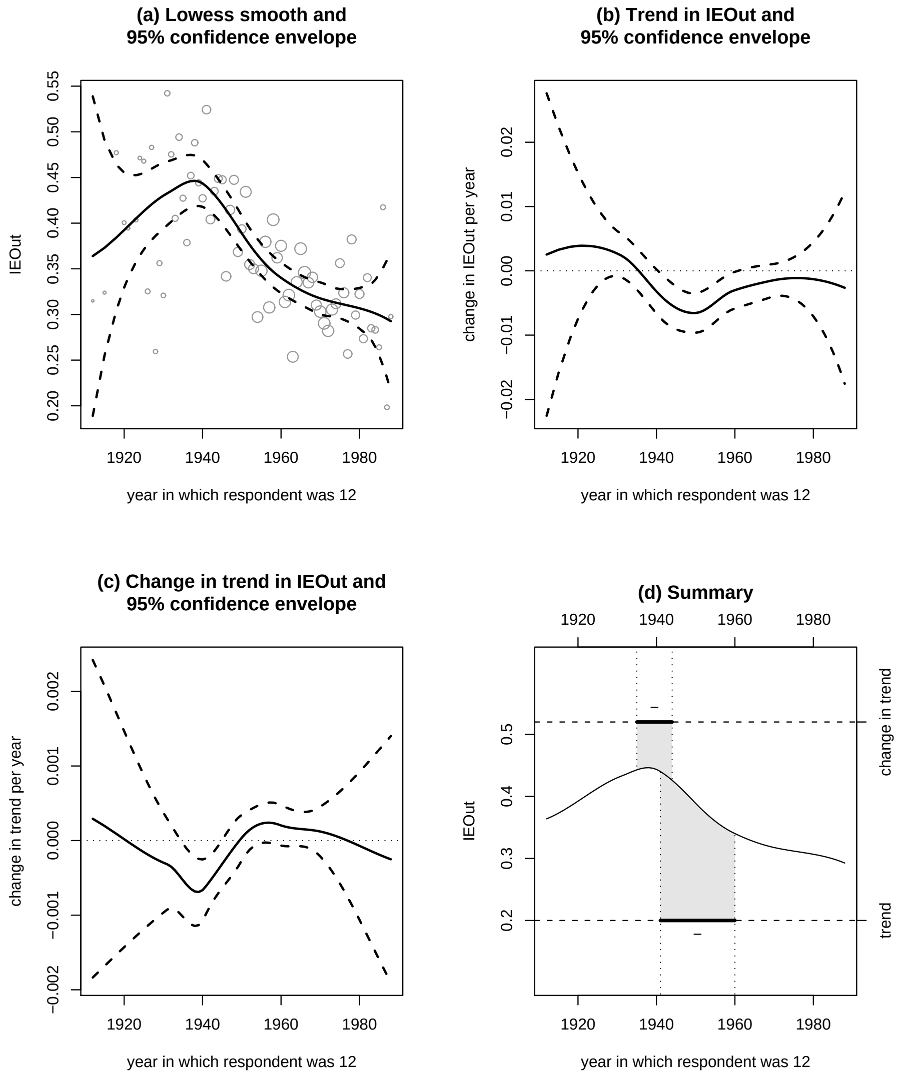

Figure 4.6: Trend in Inequality of Educational Outcomes and change in trend for women. (IEOut is measured as standardized regression coefficients. The local polynomial smooth has a span of .65 and uses weights proportional to the inverse of the variances.)

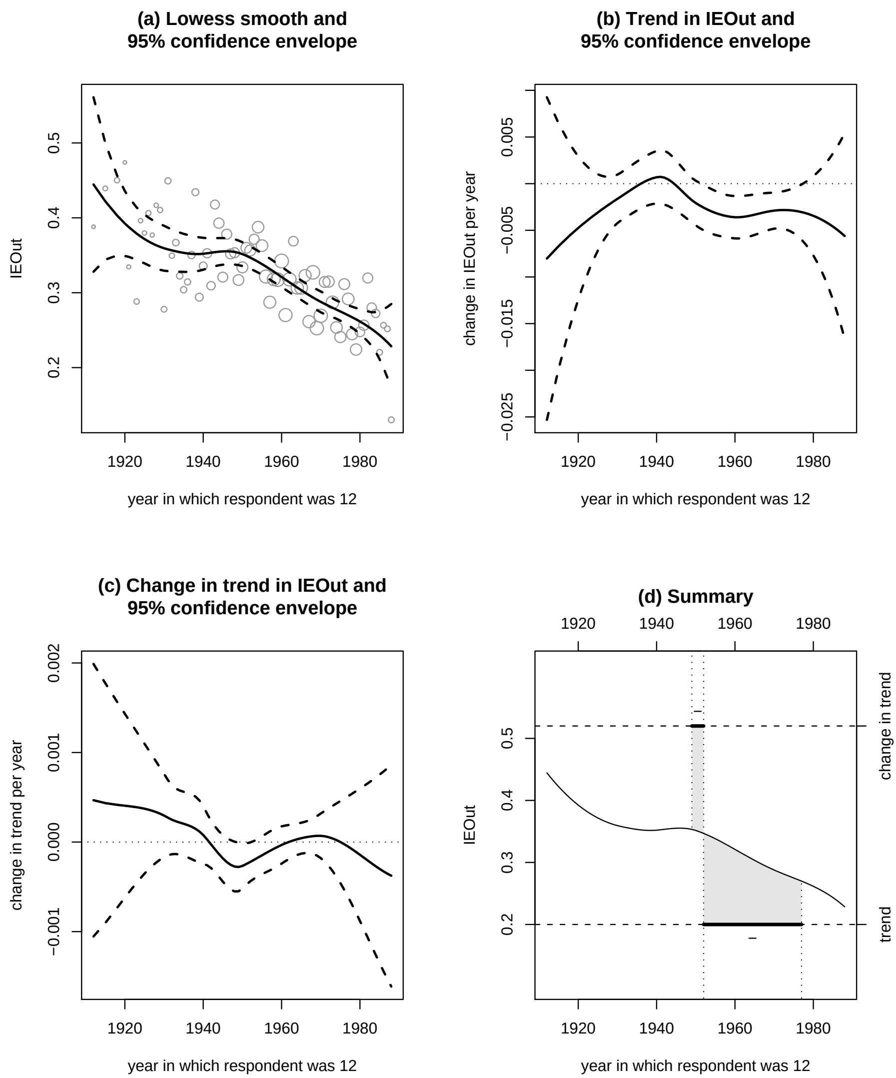

Figure 4.7: Trend in Inequality of Educational Outcomes and change in trend for men while using different scales of education and controlling for survey effects and missing data. (IEOut is measured as standardized regression coefficients. The local polynomial smooth has a span of .65 and uses weights proportional to the inverse of the variances)

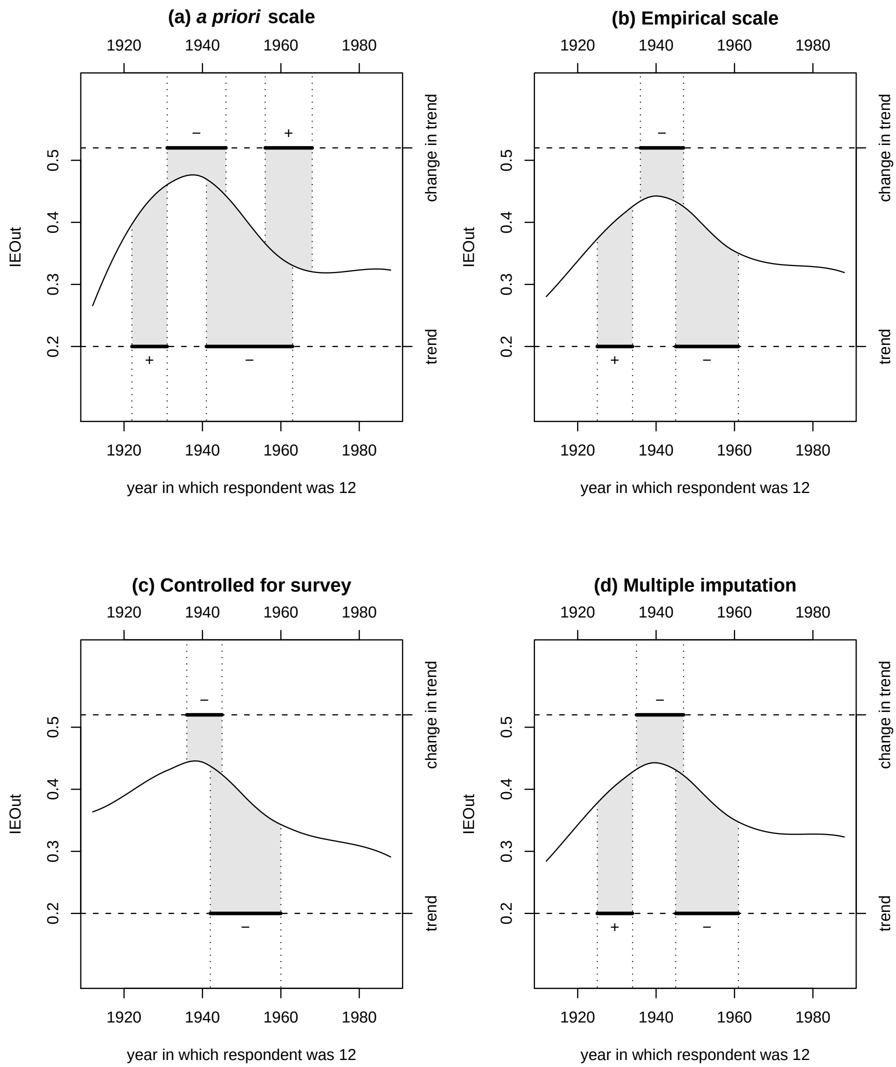

Figure 4.8: Trend in Inequality of Educational Outcomes and change in trend for women while using different scales of education and controlling for survey effects and missing data. (IEOut is measured as standardized regression coefficients. The local polynomial smooth has a span of .65 and uses weights proportional to the inverse of the variances)

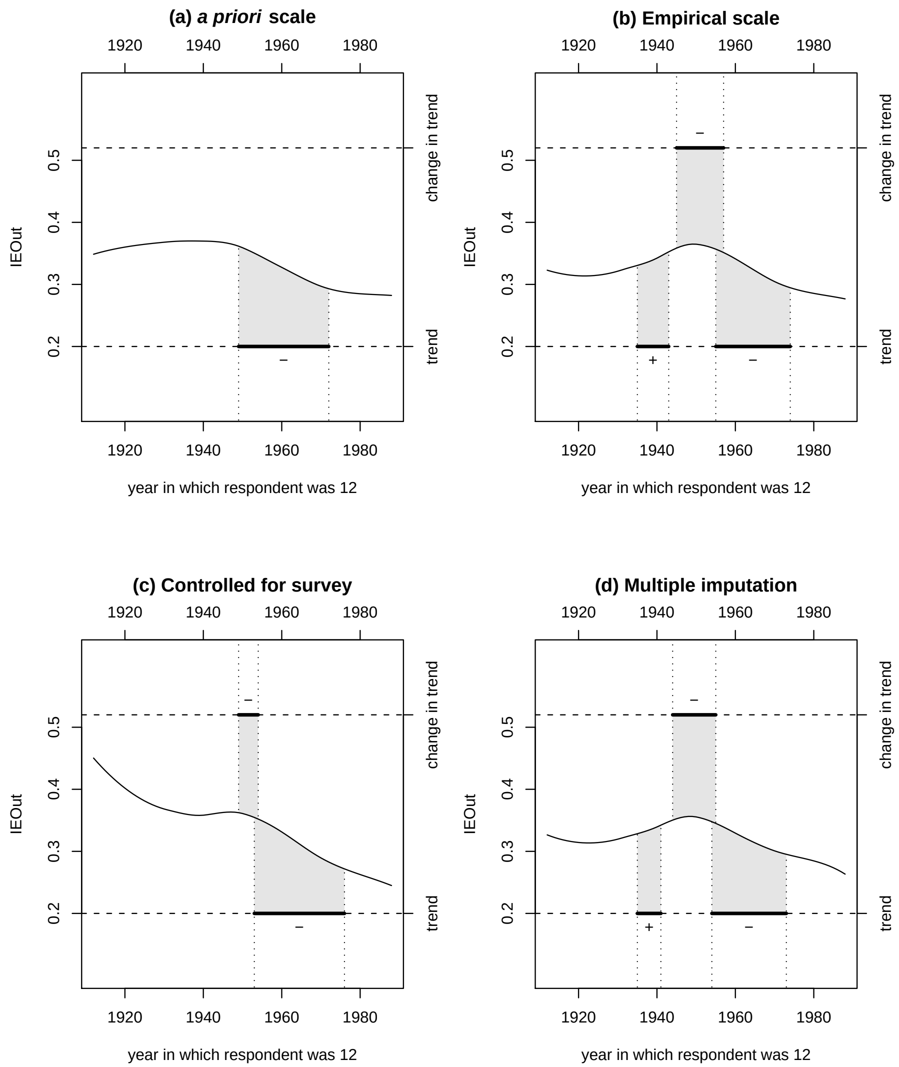

4.6 Conclusion

This chapter had a primary and a secondary aim: The primary aim was to provide a detailed description of the trend in IEOut in the Netherlands between 1912 and 1988, and in particular whether the negative trend in IEOut has decelerated. Previous studies all found a positive IEOut and an overall negative trend in IEOut, but failed to find any evidence that this trend was non-linear. This chapter did find evidence that the trend has been non-linear, but has not found the deceleration in the decreasing trend in IEOut that was expected. The results are summarized in Table 4.2, which shows periods of significant trends and changes in trends while controlling for the different potential sources of error. The most robust findings in this chapter are a period of negative trend for both men and women, which was preceded by a period of significantly accelerating trend. The presence of the period of accelerating trend indicates that previously the trend was less negative, and for men there is a solid indication that the trend was even positive. There is some evidence that the negative trend decelerated prior to becoming insignificant for men, but this deceleration is not (yet) significant. There is no indication that the negative trend for women decelerated prior to becoming insignificant, indicating that the lack of significance of the negative trend in the youngest cohorts has more to do with lack of power than with a lack of negative trend.

Table 4.2: Periods of significant trend in IEOut and change in trend in IEOut

scale	corrected for	trend		change in trend	
		positive	negative	positive	negative
men					
a priori		1922–1931	1941–1963	1956–1968	1931–1946
empirical		1925–1934	1945–1961		1936–1947
empirical	survey		1942–1960		1936–1945
empirical	missing data	1925–1934	1945–1961		1935–1947
empirical	survey and missing data		1941–1960		1935–1944
women					
a priori			1949–1972		
empirical		1935–1943	1955–1974		1945–1957
empirical	survey		1953–1976		1949–1954
empirical	missing data	1935–1941	1954–1973		1944–1955
empirical	survey and missing data		1952–1977		1949–1952

The secondary aim of this chapter was to use this analysis to investigate the degree of susceptibility of the International Stratification and Mobility File [ISMF] (Ganzeboom and Treiman, 2009) to three potential sources of error: the scaling of education, survey effects, and missing data. Controls for missing data did not change the results, but both controls for survey effects and using different scales of education did moderately influence the estimated trend. Controls for survey effects primarily influenced the older cohorts, for both men and women, while different scalings of education primarily influenced the estimated trend in older cohorts for women.

Appendix: Surveys and indicators of their quality

Table 4.3: Indicators of data quality of Dutch surveys that were post-harmonized in the International Stratification and Mobility File (Ganzeboom and Treiman, 2009)

survey number[a]	year	birth cohorts	N	response rate	% missing	# categories respondent's education	# categories father's occupation
1	1958	1912–1943	783	94	0.4	9	23
2	1967	1912–1952	1125	68	2.0	4	20
3	1967	1939–1952	333	63.1	0.3	4	21
4	1970	1912–1955	1391	74	3.1	8	62
5	1971	1912–1956	1313	76	4.4	8	57
6	1971	1921–1956	1098		2.8	5	5
7	1974	1915–1959	739	67	6.0	12	49
8	1976	1923–1961	654	69	2.2	7	59
9	1977	1918–1962	2669	70	6.9	6	64
10	1977	1918–1962	1123	64	13.7	9	59
11	1979	1921–1964	1159	65	3.1	40	55
12	1981	1923–1966	1448	83	12.6	10	63
13	1982	1924–1967	1014	74	9.5	10	64
14	1982	1929–1967	1670		13.5	8	39
15	1982	1929–1967	590	60	7.4	16	61
16[b]	1985	1928–1970	3163	41	5.6	9	67
17	1986	1928–1971	1056	83	10.7	10	62
18	1986	1928–1971	2545	57	11.0	5	39
19	1987	1929–1972	1188	82	1.6	7	30
20	1987	1929–1972	639	60	1.2	8	57
21	1987	1929–1972	686	78	10.1	7	58
22[b]	1988	1930–1973	3482		7.2	9	66
23	1990	1932–1975	1765	48	5.9	7	65
24[b]	1990	1932–1975	3303		7.5	9	67
25	1991	1933–1976	787		47.5	6	56
26	1992	1934–1978	1579	43	4.0	10	67
27[b]	1992	1934–1977	3554		7.0	9	66
28	1992	1934–1977	1624		39.1	20	64
29	1994	1936–1979	1202	47.5	11.5	5	44
30	1994	1936–1979	845	58	3.2	9	64
31[b]	1994	1936–1979	3403		13.5	8	66
32	1995	1937–1980	1639	40	11.6	9	64
33	1995	1937–1980	1615	51.5	6.1	8	65
34	1995	1956–1980	948		7.2	9	62
35	1996	1938–1981	603	36	2.6	8	59
36	1996	1938–1981	1013	37	30.5	10	14
37[b]	1996	1938–1981	3680		7.4	8	68
38	1996	1974–1981	189	42	34.4	10	14

(Continued on next page)

Table 4.3 – continued from previous page

survey number[a]	year	birth cohorts	N	response rate	% missing	# categories respondent's education	# categories father's occupation
39	1998	1940–1983	644	31	12.7	8	55
40	1998	1940–1983	1288	50	19.6	9	59
41	1998	1940–1983	1783	48	2.4	10	69
42[b]	1998	1940–1983	4039		6.1	8	67
43	1999	1941–1984	1889	43	6.5	8	64
44	1999	1941–1984	6674	66.4	23.1	8	68
45	1999	1941–1984	1062		17.4	7	61
46	2000	1942–1985	1381	40.6	2.1	10	62
47	2000	1942–1985	863	42.7	3.3	8	62
48	2002	1944–1988	1607	67.9	12.9	13	63
49	2003	1945–1988	1750		10.7	10	63
50	2003	1944–1988	6652	37.8	2.0	10	68
51	2004	1946–1988	1317	64	8.4	13	62
52	2005	1947–1988	1313	40	8.1	8	62
53	2006	1948–1988	1224	60	9.4	13	64
54	2006	1948–1988	1450	38	5.0	8	66

[a] Number refer to the data references.

[b] These are waves in a panel. Respondents in these waves are weighted to ensure that each respondent contributes only one observation.

Chapter 5

Parents and their resources:

The relative influence of the education and occupation of both parents on the educational attainment of their offspring in the Netherlands between 1939 and 1991

5.1 Introduction

The study of the influence of social background on educational attainment involves a paradox: On the one hand, it is a good thing that parents care about their children and want to help them to attain the best possible educational outcome. On the other hand, this has an undesirable consequence, as it leads to differences in educational outcomes between children from different families that do not correspond with differences in ability, talent, or motivation of the children, because families differ in the amount of social, cultural, and economic resources they have available to help their offspring. One of the tasks of the education system is to alleviate this paradox by being a separate source of resources that can, at least partially, counteract the disadvantage faced by children from parents with fewer resources. The extent to which the education system fails in reaching this goal — that is, the inequality of access in education — has been an important research topic in social stratification and mobility research (Breen and Jonsson, 2005; Hout and DiPrete, 2006), and will also be the subject of this chapter.

In this chapter I will focus on the fact that families have multiple resources available, which are contributed by both parents. In particular, this chapter will study the relative influence of the following resources: occupational status and education of the father and the mother. This will be done by answering the following two questions: First, how important were each of these resources in the Netherlands between 1939 and 1991? Second, did the relative contributions of the education and occupational status of the father and the mother to educational attainment of the offspring change in the Netherlands between 1939 and 1991?

5.2 Parental resources and their effect on the education of the offspring

When describing these parental resources, it is useful to make a distinction between who is contributing and what is being contributed.

The most obvious comparison when describing who is contributing resources is the comparison between the father and the mother, but this may not be the most relevant comparison; other alternatives are: the parent of the same sex as the offspring versus the parent of the other sex, and the parent with the highest education or occupation versus the parent with the lowest education or occupation. Moreover, these possibilities are not mutually exclusive; for instance, the fact that the father has an effect does not preclude the highest educated parent from having an effect as well. So the background variables will be entered in such a way as to allow all these combinations, in a way similar to that used by Korupp et al. (2002).

These different ways in which parents can influence the educational attainment of their offspring correspond to different hypotheses about which parent matters. The first hypothesis is based on what Goldthorpe (1983) called the 'conventional view', which states that the family's class position is determined by the father alone, because of the conventional role model in which the father is in gainful employment and the mother takes care of the children. However, this reasoning can also be used to predict the opposite: the mother's characteristics are more important for the children's educational attainment, because in this view the children are likely to interact more with the mother. Finally, one may argue that it is the resources that one brings into the household that counts, and not whether the person who brings it into the household is male or female, in which case one would expect the effect of the father's and the mother's characteristics to be equal. The second hypothesis is based on what is sometimes called the 'dominance model' (Erikson, 1984), which postulates that it is the parent with the highest status that determines the family's class position. The justification of this model can be based on the 'power model' by McDonald (1977), which assumes that these differences in status represent differences in power within the family, and that children would be influenced by the most powerful parent. However, this type of reasoning can also be turned around to come to the opposite prediction. In this view, power is at least in part derived from the occupational status, and time spent attaining occupational status competes with time spent raising children. So, it is likely that the least powerful parent spends the most time with the children, and thus would have the strongest influence. The third set of hypotheses is based on the sex-role model, which assumes that daughters are primarily oriented towards their mother and sons towards their father because the same-sex parent is perceived by the children to have more relevant information for their situation (Acock and Yang, 1984; Boyd, 1989). In principle this hypothesis could also be reversed — with the father influencing the daughter and the mother influencing the son — but it is less clear why such an arrangement would work.

As well as who is contributing resources, this chapter will also study what is being contributed. In particular, two types of resources that each parent can contribute will be considered: the highest achieved level education of the parent, and the parent's occupational status. Special attention will be paid to families in which the mother has never been in paid employment. Not only will this study try to measure the effect of the mother being a homemaker, but also two possible compensating strategies will be investigated: the father's occupation could become more important when he is the only person in the household who brings in occupational status, while the mother's education could become more important if that is her only source of status.

Finally, this chapter will also test whether the relative contributions of these resources have changed over time. Given the rapid change of the role of women in many aspects of society, it appears likely that the the mother's resources have increased in importance relative to the father's resources. However, a stability in the relative importance of the father's and mother's resources would correspond with the remarkable resilience of differences between men and women in some other areas like the division of household tasks (for example Greenstein, 2000; Gershuny et al., 1994). As a consequence, it is unclear whether to expect changing or constant relative contributions of the father versus the mother. In the case of the comparison between the parental occupational status and the parental education, there is a clear expectation about the change in their relative contributions. Occupational status is more closely related to the economic resources available in a family than parental education, and the influence of the economic resources is expected to decline over time due to two processes (De Graaf and Ganzeboom, 1993; De Graaf et al., 2000). First, economic resources influence educational attainment of the offspring by constraining the possibilities of families with insufficient economic resources. Given the economic growth in the Netherlands during the period being studied, it is expected that fewer and fewer families are constrained in their ability to send their children to school. Second, a deficiency in economic resources can easily be redressed by public policy, through subsidising education or direct subsidies to these families, and these measures have been implemented during the period under study. A similar decline in the influence of the parental education is not expected. As a consequence, the relative contribution of parental education is expected to increase.

5.3 Data and method

5.3.1 Data

The data consists of 11 surveys[1], which collected information from respondents from the Netherlands on the highest achieved level of education of the respondents, the highest achieved level of education and occupational status of their father, and highest achieved level of education and occupational status of their mother. All these surveys have been post-harmonized by Ganzeboom and Treiman (2009) as part of the International Stratification and Mobility File, ISMF. Together, these surveys contain information on approximately 11,500 respondents. This data covers the period between 1939 till 1991, as measured by the year in which the respondent was 12 (at around this age, students in the Netherlands make the most important choice in their educational career).

The highest achieved level of education of the respondents and their fathers and mothers are measured in pseudo-years, using the scale estimated in Chapter 3. The highest achieved level of education of the father and the mother has been rescaled such that it ranges between zero and one. The occupational status of the parents was measured in terms of the International Socio-Economic Index of occupational status [ISEI] (Ganzeboom and Treiman, 2003) and have also been rescaled to range between zero and one. This way, the size of the effect of the parent's education becomes comparable with the size of the effect of the parent's occupation: both measure what happens when the parent moves from the lowest position to the highest position.

In this chapter a mother is considered to have always been a homemaker if there is no information on her occupation. The homemakers are included in the analysis by setting their occupational status to zero, and adding an indicator variable to the model indicating whether or not the mother is a homemaker. The dummy for homemaker measures how much education respondents would have gained or lost if their mother had always been a homemaker rather than having the lowest status job. An interaction between the father's occupation and the homemaker dummy is added to allow the effect of father's occupational status to change when the father is the only person in the household to bring in occupational status. An interaction between the mother's education and the homemaker dummy is also added, to allow the effect of the mother's education to change when the mother's education is her only source of status.

To capture the different ways in which both parents could influence the respondent's education, the following sets of variables are added to the model:

[1] These surveys are: net92f, net94h, net95h, net95y, net96, net96y, net98, net98f, net99, net04i, and net06i, where the codes refer to the data references.

- the education and occupation of the father and the mother
- the education and occupation of the parent with the highest education or occupational status, and the education and occupation of the parent with the lowest education or occupation. This means the reference category is the parents when both have the same level of education or occupational status. Occupational statuses are considered to be equal when they differ by less then 10 ISEI points, while education is considered equal if the parents had attained the same educational category.
- the education and occupation of the parent with the same sex as the respondent, which means that the reference category is the parent of the opposite sex as the respondent. In case of female respondents, the occupation of the same-sex parent could be homemaker, so an interaction between the sex of the respondent and the homemaker indicator variable is also part of this set of variables.

So the main effects of the education of the father and the mother represent the effects when the father and the mother have the same education, and when the respondent has the opposite sex to the parent. Similarly the main effects of the occupational status of the father and the mother are the effects when the difference in occupational status between the father and the mother is less than 10 ISEI points and when the respondent has the opposite sex to the parent. All the other education and occupation variables measure the difference in effects with these reference categories.

Time is measured by the year in which the respondent was 12. This is seen as the best approximation of when any effect occurs because it is at approximately that age that students are streamed in the different tracks, which will have major consequences for their subsequent educational career. The unit of the time variable is decades since 1940. To allow for a non-linear trend, this variable is entered in the model as restricted cubic spline (Harrell, 2001) with knots at 1950, 1970, and 1980 using the `mkspline` command in Stata (StataCorp, 2007).

5.3.2 Method

The second research question requires a special model to test whether the relative impact of the different parental resources on the offspring's education changed over time. This is done by estimating a regression with parametricaly weighted explanatory variables (Yamaguchi, 2002). This model represents the null hypothesis that the effects of the parental resources may have changed over time, but that the relative impact of each of these resources has remained constant. The method will be discussed using the following simplified example: The respondent's education (`ed`) is influenced by parental education (ped) and parental occupational status (pocc), and these effects are allowed to change over time (t), as in equation (5.1).

$$\mathsf{ed} = \beta_0 + \beta_2\mathsf{t} + (1+\beta_3\mathsf{t})\left(\gamma_1\mathsf{ped} + \gamma_2\mathsf{pocc}\right) + \varepsilon \tag{5.1}$$

According to this equation, the effect of ped is $(1+\beta_3\mathsf{t})\gamma_1$ and the effect of pocc is $(1+\beta_3\mathsf{t})\gamma_2$. So, the effects of these variables are allowed to change over time, but the relative size of these effects, $\frac{(1+\beta_3\mathsf{t})\gamma_1}{(1+\beta_3\mathsf{t})\gamma_2} = \frac{\gamma_1}{\gamma_2}$, is constrained to remain constant over time. This is a so-called proportionality constraint.

The model in equation (5.1) can be estimated with maximum likelihood if we make the standard assumption that error term (ε) is normally distributed with mean 0 and a constant variance. If these assumptions are made, the alternative hypothesis, which relaxes the proportionality constraint, would then be represented by a normal linear regression with interactions between t and pedd and t and pocc. The test of the null hypothesis that the relative impact of these resources has remained constant over time is then the likelihood ratio test comparing these two models. This is implemented in Stata (StataCorp, 2007) as the `propcnsreg` package (Buis, 2007a), which is documented in Technical Materials I.

5.4 Results

The analysis started with a test of whether the relative sizes of the influence of different parental resources have remained constant. This is done by testing the model with constant relative effects of all parental resources against a model where the effects of all resources are allowed to change separately over time and between men and women, using the likelihood ratio test[2]. This results in an χ^2 value of 51.56, with 65 degrees of freedom, leading to a p-value of 0.886, which means that the null hypothesis of a constant relative effects cannot be rejected. The resulting model is shown as model 1 in Table 5.1. Table 5.1 consists of three main panels, labeled ‘constrained’, ‘trend’, and ‘main’. The parameter estimates in the panel labeled ‘constrained’ refer to the effect of the parental resources on the respondent’s highest attained level of education for men (model 1) or men and women (model 2) from the cohort that was 12 in 1940. The panel labeled ‘trend’ displays the change in effect of the parental resource variables over time and between men and women (model 1) or only over time (model 2). The panel labeled ‘main’ captures the effects of other variables that influence educational background. This panel contains the main effects of the variables specified in the panel ‘trend’, but could also have contained other control variables.

[2]The model with the proportionality constraint is presented as model 1 in Table 5.1, while the parameter estimates of the unconstrained model are not shown due to the large number of parameters in this model.

Table 5.1: Parameter estimates of models explaining highest achieved level of education

			Model 1		Model 2	
			b	se	b	se
constrained	occupation	father	2.571	(0.52)	3.437^a	(0.26)
		mother	3.442	(0.53)	3.437^a	(0.26)
		highest	-0.013	(0.46)		
		lowest	0.124	(0.64)		
		same sex	0.477	(0.45)		
		homemaker	-0.746	(0.24)	-0.625	(0.21)
		homeXfemale	0.465	(0.23)		
		homeXfather	1.367	(0.54)	1.955	(0.44)
	education	father	2.421	(0.36)	2.470^b	(0.20)
		mother	2.133	(0.38)	2.470^b	(0.20)
		highest	1.042	(0.26)	1.246	(0.23)
		lowest	-0.983	(0.41)	-1.135	(0.41)
		same sex	0.081	(0.35)		
		homeXmother	1.006	(0.45)	0.945	(0.46)
trend	year_1		-0.144	(0.03)	-0.158	(0.02)
	year_2		0.075	(0.03)	0.078	(0.02)
	female		0.125	(0.11)		
	femaleXyear_1		-0.050	(0.06)		
	femaleXyear_2		0.017	(0.05)		
	constant		1.000	.	1.000	.
main	year_1		0.617	(0.14)	0.681	(0.12)
	year_2		-0.429	(0.15)	-0.437	(0.12)
	female		-2.148	(0.40)	-1.684	(0.23)
	femaleXyear_1		0.576	(0.21)	0.415	(0.12)
	femaleXyear_2		-0.099	(0.22)	-0.074	(0.12)
	constant		7.945	(0.29)	7.790	(0.25)
log likelihood			-29951.4		-29959.2	

^{a,b} entries with the same superscript are constrained to be equal.

Table 5.2: Constraints on the effects of the parental resources (Wald tests)

Null hypothesis	occupation			education		
	χ^2	df	p	χ^2	df	p
female = 0	24.88	1	0.000	31.97	1	0.000
father = mother	1.64	1	0.201	0.43	1	0.511
highest = same = lowest	0.11	2	0.947	16.22	2	0.000
same sex = different sex	4.24	2	0.120	0.05	1	0.830

The analysis continued with a description of the effects of the parental resources. These effects are shown in the panel labeled 'constrained' in Table 5.1. This description can be split into two parts. The first part has to do with which parent contributes the resource: only the father, the father and the mother, the parent with the highest and the lowest occupational status or education, and/or the parent with the same and the opposite sex. Model 1 simultaneously allows all these effects. These effects were tested and these tests are reported in Table 5.2. The first row in this table reports the test that only the father contributes, this is the conventional hypothesis. This hypothesis is rejected for both the parental education and the parental occupational status. The second row tests whether there is a difference in effect between the occupational status and the education of the father and the occupational status and education of the mother. The hypothesis that the effects are the same for both fathers and mothers cannot be rejected for parents' occupation nor for parents' education. The third row tests the dominance hypothesis: whether the effect of the parent with the highest education or occupational status differs from the effects when both parents have the same occupational status or education, and whether the effects of the parents when both parents have the same education or occupational status differs from the effect of the parent with the lowest education or occupational status. The hypothesis that these effects are the same must be rejected for the education of the parents, but this is not the case for the parents' occupational status, indicating support for the dominance hypothesis for parental education but not for parental occupational status. Finally, the last row tests the sex role hypothesis: whether the effect of the mother on the daughter and the father on the son is different from the effect of the mother on the son and the father on the daughter[3]. The hypothesis of no difference in effect of the parent with the same sex as the respondent and the parent with the opposite sex as the respondent could not be rejected, neither for the effect of parental education nor for the effect of parental occupational status. This provides evidence against the sex-role hypothesis.

[3]Notice that the effect of the occupation of the parent with the same sex as the respondent is captured by two variables, the occupational status of the same sex parent and the interaction between homemaker and female. So this is a 2 degree of freedom test for occupation and a 1 degree of freedom test for education.

The second part of the description of the effect of parental resources has to do with which resource contributes most. Two types of resources have been distinguished: the occupational status of the parent, and the education of the parents. Of particular interest in this case are the parameters of father and mother in the first panel, which represents the effects of the father's and the mother's occupation or education in 1940 when they have the same occupational status or education as their partner, are not of the same sex as the sex of their offspring, and the mother has worked. It indicates that the effect of parental occupational status is stronger than the effect of parental education. Within model 1 this difference is not significant ($\chi^2(2) = 3.08$, p = .214), but within model 2 parental occupational status has a significantly stronger effect than parental education ($\chi^2(1) = 8.54$, p = .004).

Model 1 can be further simplified by forcing the effects of the resources to be the same for male and female respondents, that is constraining the effects of female, femaleXyear$_1$, and femaleXyear$_2$ in the second panel of Table 5.1 to be zero. All these constraints together result in the simplified model 2 in Table 5.1. The parameters can be interpreted in the following way: Within the sub-panel labeled 'occupation', the parameters for father and mother are the effects of the father's and mother's occupational status on the respondent's education in 1940 if the mother has not always been a homemaker. It shows that if a parent moves from the lowest to highest status occupation, the education of the offspring is expected to increase by 3.5 pseudo-years. The effect of the variable homemaker indicates the difference in pseudo-years of education between respondents whose mother has always been a homemaker and whose mother had a job with the lowest status. So the offspring is likely to attain more education when the mother has had the lowest status job as apposed to being a homemaker. The effect of homeXfather shows that when the mother has always been a homemaker, the father's occupational status increases by about 2.0 pseudo-years. This means that the negative effect of the mother being a homemaker can be decreased or even reversed by an increase in the father's occupational status. The sub-panel labeled 'education' shows that increasing a parent's education from the lowest to the highest level would result in an increase in the offspring's education of 2.5 pseudo-years if the father and the mother have the same education, and that this effect increases by 1.2 pseudo-years if the parent is the highest educated parent, and decreases by 1.1 pseudo-years if the parent is the lowest educated parent. The effect of the interaction term homeXmother shows that if the mother has always been a homemaker, the effect of her education increases by about a pseudo-year. As a consequence, the effect of the mother being a homemaker can become less negative or even positive when the mother has a higher level of education.

These effects are also represented in Figure 5.1, together with how they changed over time. Due to the proportionality constraint, the shape of the trend is the same for

Figure 5.1: Effects of parental resources on respondent's education

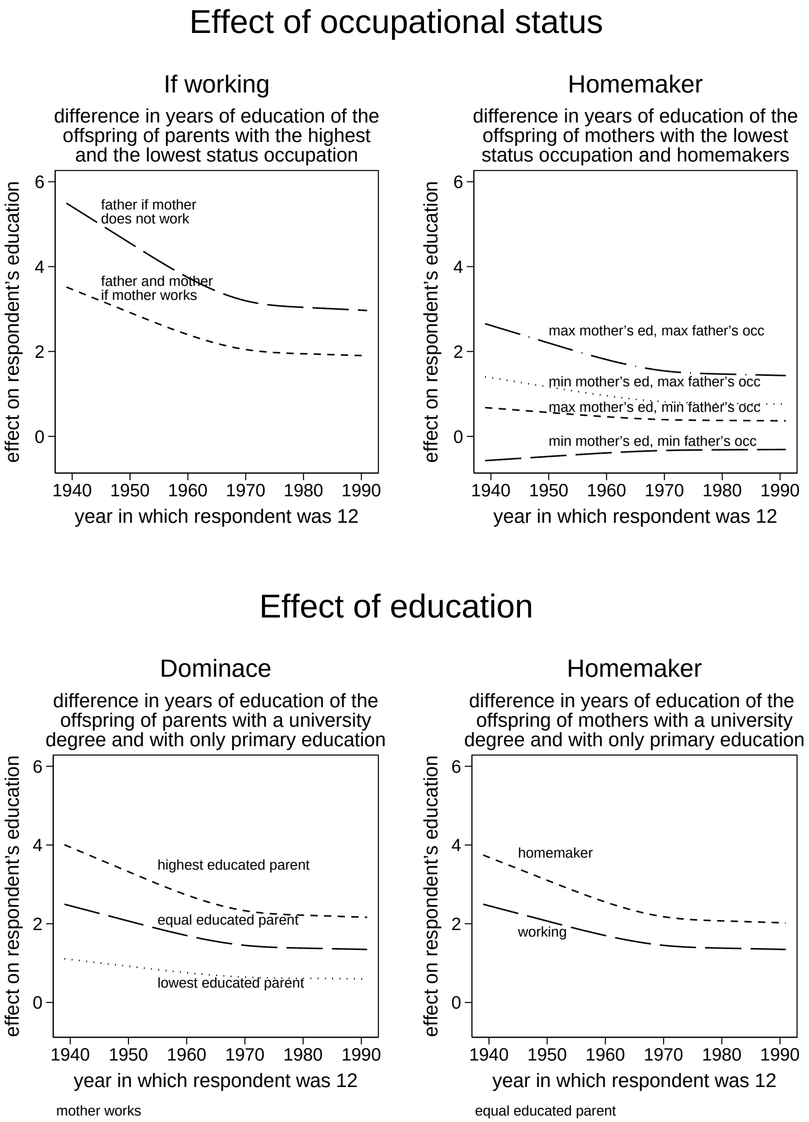

all family background variables. It shows that the effects decrease over time, but that this decrease slows down. The time trend is in Table 5.1 represented by the restricted cubic spline terms year_1 and year_2, which were parameterized in such a way that if year_2 is not significant, the trend is not significantly different from a linear trend, so Table 5.1 shows that this slowing down of the trend is statistically significant.

5.5 Conclusion

This chapter started with the notion that parents have multiple resources available with which they can help their offspring. This chapter focussed on two of these: parental education and parental occupational status. Two questions were asked about this: First, how important are each of these resources in the Netherlands between 1939 and 1991? Second, did the relative contributions of the education and occupational status of the father and the mother to educational attainment of the offspring change in the Netherlands between 1939 and 1991?

The first question was split up into two parts:

1. which parent contributes most to the educational attainment of the offspring:
    - the father or the mother, or
    - the parent with highest or lowest education or occupation, or
    - the parent with the same sex as the respondent or the opposite sex, or
    - any combination of these three?
2. what parental resource contributes most to the educational attainment of the offspring: their education or occupational status?

The analysis showed that as long as the mother works, it does not matter who brings in the resources. The only exception is that the education of the highest educated parent has a larger effect than the effect of education if both parents have the same level of education, which in turn is larger than the effect of the lowest educated parent. Otherwise, the effects of the father's characteristics are the same as the effects of the mother's characteristics, there is no difference in the effects of the education and occupational status of the parent with the same sex as the respondent and the parent with the opposite sex to the respondent, and there is no difference in the effects of the parent with the highest, same, and lowest occupational status. Having a mother who has always been a homemaker decreases the respondent's expected level of education compared to respondents from mothers with the lowest status job. However, it also increases the effects of father's occupational status and of mother's education.

The negative effect of the mother being a homemaker on the offspring's education becomes positive when the mother is highly educated and/or the father has a high status job. The parent's occupational status appears to have a stronger influence than the parent's education. This could be due to how parental education and occupation were standardized. Both were standardized such that their effect represents the effect of moving from a parent with the lowest education/occupational status to a parent with the highest education/occupational status. Because there are only a limited number of educational categories, the distribution of education is more restricted than the distribution of occupational status. As a consequence, the difference between the highest and lowest educational category is likely to be smaller than the difference between the highest and lowest occupational status. The fact that the unit of education implies a smaller step than the unit in occupational status could (in part) explain the difference in effect.

The expected answer to the second question was that over time the resources of the mother could have become more important due to the changing role of women in Dutch society during this period. In addition, the impact of occupational status was expected to decline because occupational status was expected to be more closely related to economic resources, and economic growth and government policy meant that lack of economic resources in a family has become less of a constraint for attaining education. However, no such changes were found in this study. A possible reason for this could be lack of statistical power. The test of this hypothesis was a test that the effects of all the resources on the offspring's education changed over time in such a way that the relative differences in effect remained constant, which is a proportionality constraint. This is a rather subtle constraint, and a test of this constraint is thus a test with a rather low statistical power.

The two main findings of this chapter are that it matters relatively little which parent brings in the resources as long as the mother works, and that no evidence was found that the relative contributions of different family resources have changed over time. The lack of evidence for a change in the relative contributions was not expected, but it has a fortuitous practical consequence for social stratification and mobility research: a significant part of this literature has used only a single indicator of parental resources to estimate the effect of family background on educational attainment of the offspring, most commonly the father's occupational status. A negative trend in the effect of father's occupational status would in that case be open to a number of interpretations: either the educational system has become more open to people from different backgrounds, or father's occupational status has become an increasingly bad proxy for family background as fathers have lost influence relative to mothers, or father's occupational status may have become less important but other family background characteristics, like education, may have remained constant or even increased in importance. However, the first interpretation seems to be the correct one, as no changes in the relative effects have been found. So, the use of a single indicator for family background is still a reasonable strategy, especially when only one indicator is present in the data.

Chapter 6

Not all transitions are equal:

The relationship between inequality of educational opportunities and inequality of educational outcomes

6.1 Introduction

Social stratification research has long been concerned with the relationship between family socioeconomic status (SES) and offspring's educational attainment (Breen and Jonsson, 2005; Hout and DiPrete, 2006). A strong positive association between the two implies that respondents with higher SES backgrounds are more likely to achieve higher levels of education than respondents with lower SES backgrounds. For this reason, the strength of the relationship is often termed 'Inequality of Educational Opportunity' or IEO (Boudon, 1974; Mare, 1981). IEO can be measured in a variety of ways, and these different measures tend to lead to seemingly different conclusions. This chapter will focus on two of these measures of IEO: the association between family SES and the highest achieved level of education, and the association between family SES and probabilities of passing from one educational programme to the next. These will be called Inequality of Educational Outcome (IEOut) and Inequality of Educational Opportunity proper (IEOpp) respectively, while IEO will be used as a generic term. IEOut focusses on the end result of the educational process, which is often of interest as this result, the highest achieved level of education, is the most visible result of education in subsequent areas of life like the labor market or the marriage market. IEOpp focusses on the process of attaining education. Attaining a level of education is something that typically happens over a long period of time and is usually split up into different steps, for example finishing primary eduction, finishing secondary education, etc. Knowing the influence of SES at each of these transitions can give a more complete picture of how IEO came about. So, these two measures of IEO capture different aspects of IEO: IEOut describes inequality of the outcome of the process of attaining education, while IEOpp describes inequalities in that process itself. The aim of this chapter is to show how estimates of IEOpp and IEOut can complement one another. The key challenge when dealing with complementary models is to find a way to move beyond just presenting separate results from different models to an integrated discussion of the results that shows how the different results are related to one another.

This is done by demonstrating that there is a relationship between IEOpp and IEOut in the form of a decomposition of IEOut as a weighted sum IEOpps. This means that the IEOpps (the process) lead to IEOut (the outcome), but that not every IEOpp (that is, every step in the process) is of equal importance for achieving the outcome. Moreover, as will be shown below, the importance of each IEOpp for the IEOut can differ across groups. A clear example of this is the differences in the importance of the transition between primary education and secondary education between cohorts. In most industrialized countries virtually all students within the recent cohorts remain in education after the primary level. As a result, any inequality at this first transition only affects a few (or no) students, and is thus not very important for IEOut. The situation was quite different at the beginning of the twentieth century: at that time many more students failed to continue after primary education, so the IEOpp for the transition between primary and secondary education was much more important for the IEOut than it is now. Within the decomposition developed in this chapter there will be two additional reasons why the importance of a transition can differ across groups: the importance of a transition will increase as the proportion of people at risk increases, and when the difference in the value of the expected highest attained level of education between those that pass and those that fail increases. All three are substantively interpretable ways in which the distribution of education — that is, for each educational programme the proportion of people that has that program as their highest achieved level of eduction — can influence IEOut. This decomposition thus leads one to relate IEOpp and IEOut to one another as two complementary descriptions of IEO, and allows one to investigate the effect of changes in the distribution of education on IEOut. The fact that IEOut and IEOpp are related is not new, Mare (1981) already established that, but the use of this relationship to create an integrated analysis of IEOpp and IEOut and to study the impact of educational expansion on IEOut is new to the best of my knowledge.

This chapter will begin with a description of a number of models of educational inequality. This will be followed by a discussion of the model proposed by Mare (1981), and the derivation of the relationship between IEOpp, IEOut, and the distribution of education. In the next section the decomposition will be illustrated by applying it to differences in IEOut between men and women and across cohorts that were 12 years old in the Netherlands between 1905 and 1991.

6.2 Different models of IEO

A variety of different models have been proposed and used for studying IEO. These different models tend to emphasize different aspects of IEO. For example much of the early research focusses on inequality in the end result by studying the association between family background and highest achieved level of education (Blau and Duncan, 1967; Duncan, 1967; Hauser and Featherman, 1976). This research was supplemented by Boudon (1974) and Mare (1980, 1981), who studied educational inequality during the process of attaining education as the effect of family background on the probability of passing steps between educational programmes. In particular, Mare (1980, 1981) proposed the use of the sequential logit model for estimating IEOpps. Estimates of IEOpp and IEOut are now often treated as competing representations of educational inequality. The reason for that is that Mare (1981) showed that there is an relationship between IEOpp and IEOut which involved the transition probabilities, but presented this relationship as a black box. The main point he made was that differences in these estimates of IEOut between cohorts are in part due to differences in the distribution of education. These effects can be considerable, since the distribution of education varies substantially over cohorts. In almost all countries, people born in later cohorts have attained more education, a process that has been termed 'educational expansion' (Hout and DiPrete, 2006). Furthermore, Mare (1981) showed the IEOpps control for this effect of educational expansion. This led Mare (1981) to argue that the IEOpps are a more 'pure' measure of IEO. Since then, the literature has approached the relationship between IEOut, the IEOpps, and the distribution of education as a black box.

This practice leads one to ignore two opportunities. First, the complementary nature of the information contained in estimates of IEOpps and IEOut are not fully used when treating the relation between these two as a black box. IEOpp and IEOut are natural complements as the former describes the process of attaining education while the latter describes the outcome of that process. Some studies report both estimates for the IEOpps and the IEOut, (for example Shavit and Blossfeld, 1993) but these do not relate the two types of estimates to one another. Second, this practice makes it hard to study the impact of educational expansion on IEO, because one explicitly controls for changes in the distribution of education. Those studies that have investigated the relationship (Mare, 1981; Smith and Cheung, 1986; Nieuwbeerta and Rijken, 1996) compare the observed IEOut with the simulated results of two counterfactual scenarios, those being that either the distribution of education remained unchanged and IEOpp changes as observed; or that the distribution of education changes as observed, but IEOpp remains unchanged. Simulations such as these can tell us how much IEOut is affected by changes in the distribution of education and changes in IEOpps, but do not offer us any insights as to why. This leads to the following two questions:

> How are IEOut and IEOpp related to one another, and how can this relation be used for a meaningfully integrated analysis of IEOpp and IEOut?
>
> How are IEOut and the distribution of education related to one another, and how can this relation be used for an analysis of the influence of changes in the distribution of education on IEOut?

These questions are answered by showing that the standard model for estimating IEOpps, the sequential logit model proposed by Mare (1981), implies an estimate of IEOut, which can be decomposed into a weighted sum of the IEOpps. Moreover, it will be shown that each IEOpp's weight depends on the distribution of education in three substantively interesting ways. An IEOpp receives more weight if 1) the proportion of people 'at risk' of making that transition increases; 2) the proportion passing that transition is closer to 50%, that is, passing or failing that transition cannot be regarded as almost universal; and 3) the difference in expected level of education between those who pass and those who fail to make the transition increases, that is, the expected gain from passing increases. This decomposition of IEOut into a weighted sum of IEOpps provides a link between IEOpp and IEOut and a way of conducting an integrated analysis of the two. The decomposition of the weights into the product of its three elements provides a link between the distribution of education and IEOut and a way of showing the influence of changes in the distribution of education on IEOut. The decomposition of IEOut into IEOpps and weights has been implemented in Stata (StataCorp, 2007) in the `seqlogit` package (Buis, 2007b), which is documented in Technical Materials II.

This decomposition does not require a new model, it is just a different way of presenting the results of a sequential logit model. This means that the critique by Cameron and Heckman (1998) on the sequential logit model also applies to this decomposition. Their argument starts with the observation that it is very likely that not all variables that influence the probability of passing a transition are observed. In this case the sequential logit model will estimate the effect of the observed explanatory variables on the proportion of respondents that pass a transition averaged over these unobserved variables rather than on an individual's probability of passing the transitions. The problem is that the group level effects measured by the sequential logit model will not be the same as the individual level effects, even if the unobserved variables are non-confounding variables. The easiest solution is to interpret the results of the sequential logit model as a description of differences between different groups rather than interpret the results as individual-level effects. Alternatively, one can try to adapt the model to take unobserved heterogeneity into account. This is obviously a difficult problem, as one tries to control for variables that have not been observed, and a consensus on the best way of doing this has yet to appear. A discussion of the various solutions proposed to solve this problem is beyond the scope of this chapter, so the main focus of this chapter will be on the effects on group-level transition rates rather than individual-level effects. However, the decomposition can be applied to some of the models that have been proposed for estimating individual-level effects (for example: Mare 1993 and Chapter 7 of this dissertation), and generalizations of the decomposition for these models will be briefly discussed.

6.3 The relationship between inequality of educational opportunities and outcomes

In this section I will derive and discuss a decomposition of an estimate of IEOut into a weighted sum of IEOpps. This decomposition starts with the model for IEOpps proposed by Mare (1981), which I will refer to as the sequential logit model (following Tutz (1991)). This model is also known under a variety of other names: sequential response model (Maddala, 1983), continuation ratio logit (Agresti, 2002), model for nested dichotomies (Fox, 1997), and simply the Mare model (Shavit and Blossfeld, 1993). Consider, for instance, a hypothetical education system consisting of four levels: no education, primary education, secondary education, and tertiary education as represented in Figure 6.1. Figure 6.1 shows how respondents face three transitions in this system: they can attend primary education or opt for no education at all; if they opt for primary education they can choose to leave the system once they have completed primary education, or go on to secondary education; and if they opt for secondary education, they can then either choose to leave once they have completed this level or go on to tertiary education. The implication is that if someone's highest-achieved level of education is primary education, then that person was 'at risk' of passing the first two transitions, but not the third. Furthermore, it implies that the person passed the first transition, but failed the second.

The model assumes that one has to be 'at risk' of passing a transition — that is, to have passed through all lower transitions — in order to make a decision at that transition about whether to continue in education or to leave the system. Aside from this, these decisions are assumed to be completely independent. As a result, one can estimate the IEOpp by running separate logistic regressions for each transition on the appropriate sub-sample (Mare, 1980). This model is shown in equation (6.1).

Figure 6.1: Hypothetical educational system

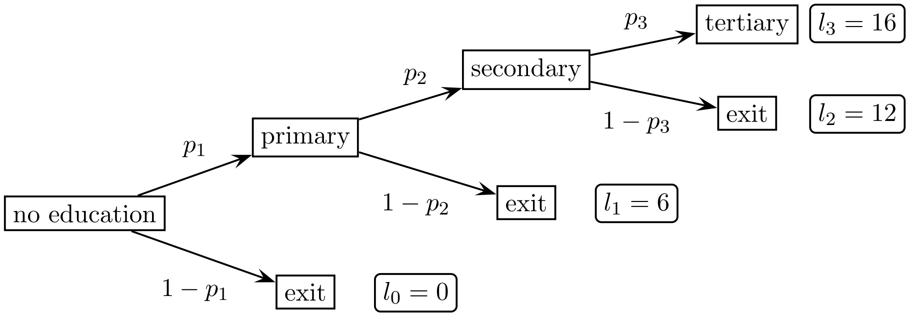

$$
\begin{aligned}
\hat{p}_{1i} &= \frac{\exp(\alpha_1 + \lambda_1 SES_i + \beta_1 x_i)}{1 + \exp(\alpha_1 + \lambda_1 SES_i + \beta_1 x_i)} \\
\hat{p}_{2i} &= \frac{\exp(\alpha_2 + \lambda_2 SES_i + \beta_2 x_i)}{1 + \exp(\alpha_2 + \lambda_2 SES_i + \beta_2 x_i)} \quad \text{if} \quad pass_{1\,i} = 1 \\
\hat{p}_{3i} &= \frac{\exp(\alpha_3 + \lambda_3 SES_i + \beta_3 x_i)}{1 + \exp(\alpha_3 + \lambda_3 SES_i + \beta_3 x_i)} \quad \text{if} \quad pass_{2\,i} = 1
\end{aligned} \tag{6.1}
$$

The probability that person i passes transition k is $\hat{p}_{ki}$. The IEOpp belonging to transition k is λ_k, the constant for transition k is α_k, and the effect of a control variable x_i is represented by β_k. Whether or not individual i has passed the previous transition is indicated by the indicator variable $pass_{k-1\,i}$. It is assumed that everybody is at risk of passing the first transition. The differences in IEOpp between men, women, and cohorts can be obtained by adding the appropriate interaction terms to the model.

In order to make a link between the IEOpps (the λ_ks) and IEOut, it is necessary to assign a value (l_k) to each level of education. By assigning values to each educational level, it becomes possible to use the sequential logit model to calculate the expected highest achieved level of education ($E(L_i)$). The results from the sequential logit are used to compute predicted probabilities for passing each transition, and the expected highest achieved level of education is the sum of the value of each level of education times the probability of attaining that level. This is set out in equation (6.2). The probabilities and values assigned to each level can be derived from Figure 6.1[1].

[1]The values that are assigned to each of the levels in Figure 6.1 are typical for when these values are based on years or pseudo-years of education, but this decomposition is not limited to this metric.

$$E(L_i) = (1 - \hat{p}_{1i})l_0 + \hat{p}_{1i}(1 - \hat{p}_{2i})l_1 + \hat{p}_{1i}\hat{p}_{2i}(1 - \hat{p}_{3i})l_2 + \hat{p}_{1i}\hat{p}_{2i}\hat{p}_{3i}l_3 \qquad (6.2)$$

The family's SES is part of equation (6.2) through the $\hat{p}_{ki}$s described in equation (6.1). Equation (6.2) can be understood as a regression equation showing a non-linear relationship between a family's SES and the highest achieved level of education. Using a sequential logit model to derive such a (non-linear) regression is unusual. A more common method for estimating IEOut is to use a linear regression of highest achieved level of education on family SES (for example, Blau and Duncan, 1967; Shavit and Blossfeld, 1993). The advantage of the non-linear model derived from the sequential logit model over the linear model is that the non-linear model provides the link between the IEOpps and the IEOut. Moreover, the non-linear model takes the bounded nature of the dependent variable into account, as it can never lead to predictions below the lowest level of education or above the highest level of education.

Recall that IEOut is the effect of a family's SES on the respondent's expected highest achieved level of education, or, in other words, how much the expected highest achieved level of education changes if a family's SES changes[2]. Consequently, IEOut is the first derivative of equation (6.2) with respect to a family's SES. This derivative is shown in equation (6.3). A step-by-step derivation is set out in the appendix to this chapter.

$$\begin{array}{llllllll}
\frac{\partial E(L_i)}{\partial SES} & = & & & & & & \\
\{1 & \times & \hat{p}_{1i}(1-\hat{p}_{1i}) & \times & [(1-\hat{p}_{2i})l_1 + \hat{p}_{2i}(1-\hat{p}_{3i})l_2 + & & & \\
 & & & & \hat{p}_{2i}\hat{p}_{3i}l_3 - l_0] & \} & \lambda_1 & + \\
\{\hat{p}_{1i} & \times & \hat{p}_{2i}(1-\hat{p}_{2i}) & \times & [(1-\hat{p}_{3i})l_2 + \hat{p}_{3i}l_3 - l_1] & \} & \lambda_2 & + \\
\{\hat{p}_{1i}\hat{p}_{2i} & \times & \hat{p}_{3i}(1-\hat{p}_{3i}) & \times & [(l_3 - l_2)] & \} & \lambda_3 &
\end{array} \qquad (6.3)$$

Equation (6.3) shows that IEOut $\left(\frac{\partial E(L_i)}{\partial SES}\right)$ is a weighted sum of the IEOpps (the λ_ks). The weights (the sections between curly brackets) consist of three parts, all of which are related to the distribution of education. These are:

1. The predicted proportion of people at risk of passing a transition. For the first transition, this proportion is 1; for the second it is the proportion of students who complete primary education, $\hat{p}_{1i}$; and for the third transition, it is the proportion who completed secondary education, $\hat{p}_{1i}\hat{p}_{2i}$. Substantively, this means that a transition is more important when more people are at risk of passing it.

[2]More precisely, the measure of IEOut used in this chapter studies how the average highest achieved level of education of a group of respondents with the same family SES reacts to a change in the family SES rather than an individual-level effect, as was discussed before.

2. The variance of the indicator variable showing who passed and who failed the transition, $\hat{p}_{ki}(1 - \hat{p}_{ki})$. This variance is a function of the predicted probability of passing. This is lowest if virtually everybody passes or fails, and is highest when the probability of passing is .5. This makes sense at a substantive level, because if only a few people pass or fail a transition, then any inequality at this stage will only affect a few people.

3. The differences between the expected level of education of those who pass the transitions and those who do not. These are the parts in the square brackets. For instance, the expected level of education of those who pass the first transition is $(1 - \hat{p}_{2i})l_1 + \hat{p}_{2i}(1 - \hat{p}_{3i})l_2 + \hat{p}_{2i}\hat{p}_{3i}l_3$ and the expected level of education for those that fail the first transition is l_0. The difference between the two is the expected gain from passing the first transition. The substantive interpretation of this is that a transition becomes more important if passing it leads to a greater expected increase in the highest achieved level of education.

The result is summarized below. IEOut is a weighted sum of IEOpps, and the weights are the product of the proportion at risk, the variance, and the expected gain in level of education resulting from passing.

$$\begin{aligned} \text{IEOut}_i &= \sum_{k=1}^{K} (\text{weight}_{ki} \times \text{IEOpp}_k) \\ \text{weight}_{ki} &= \text{at risk}_{ki} \times \text{ variance}_{ki} \times \text{ gain}_{ki} \end{aligned}$$

Each respondent will have its own IEOut and set of weights because the weights are based on the predicted probabilities of passing the transitions, and these probabilities will differ between persons depending on their values on the explanatory variables. In this chapter this decomposition will be summarized by computing the decomposition for an individual with average values on the explanatory variables. This is not the only way one can summarize the IEOuts. For example, one can compute the IEOut for each individual and average those. This 'averaged IEOut' can also be decomposed into a weighted sum of IEOpps, where the weights are now the average of the weights predicted for each individual. However, these averaged weights can no longer be decomposed as the product of its three constitutive elements[3]. This is why the IEOut of a person with average values on its explanatory variables is preferred over the 'averaged IEOut'.

[3] The reason for this is that the weight is a product of variables, and the average of a product of variables is not the same as the product of the averages of these variables.

As was discussed before, this decomposition is just a different way of representing the results from a sequential logit model, so the criticism by Cameron and Heckman (1998) also applies here. However, this decomposition can be extended to models that estimate individual-level IEOpps as long as the individual-level IEOpps are estimated by modelling the transition probabilities using a logistic curve, as is the case in (Mare, 1993) and Chapter 7. In both articles, certain assumptions are made concerning the distribution of the unobserved variables, and the IEOpps are estimated given these assumptions. The presence of the unobserved variables complicates the estimation in ways that are beyond the scope of this chapter, but within the context of equations (6.1), (6.2), and (6.3) the unobserved variable is not different from the observed variables. In this case one can create predicted probabilities for someone with average values on both the observed and unobserved variables and use those to compute the decomposition in equation (6.3).

In summary, the main advantage of the decomposition proposed in this chapter is that it allows for an integrated discussion of IEOpps and IEOut and a way of studying the influence of changes in the distribution of education on IEOut. This makes it possible to make full use of the complementary nature of IEOpp and IEOut, and to study the influence of factors such as educational expansion on IEOut. One can easily extend this argument, allowing us to study the roles played by gender educational inequality, racial educational inequality, or differences in the distribution of education between countries. A graphical representation of this decomposition is presented during the empirical discussion.

6.4 Empirical application

This section will illustrate how the relationship between IEOpp, IEOut, and the distribution of education can be used to gain a more complete picture of IEO. In particular, this section will describe the relationship between IEOpp and IEOut and the influence of educational expansion and gender inequality on IEOut in the Netherlands for cohorts that were 12 years old between 1905 and 1991.

6.4.1 The Dutch education system

The aim is to estimate a sequential logit model for the Netherlands and use the results to compute the decomposition of IEOut into IEOpps and their weights. The challenge is to come up with a model for the Dutch education system that provides a good representation of the education system during the entire period under study and where the assumption that each level can be achieved via only one route through the education system is plausible. The strategy used for meeting these challenges is to create a

Figure 6.2: Simplified model of the Dutch education system

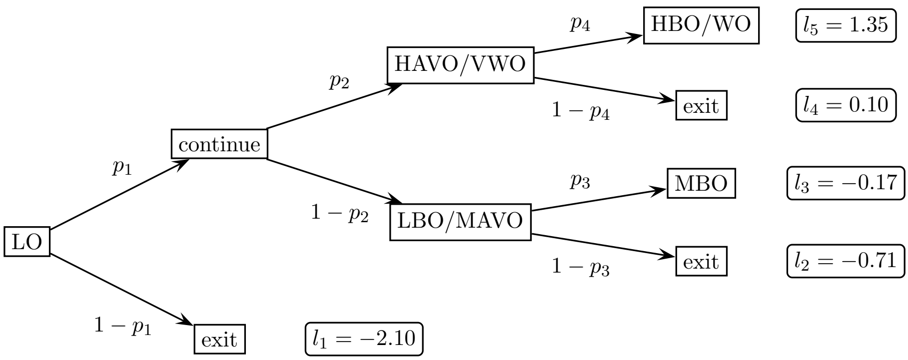

stylized model of the Dutch education system by combining educational programmes into 'rougher' categories. This helps with keeping the model representative for the entire period, because even though the position of individual educational programmes within the Dutch education system could have changed over time, the positions of the rougher categorizations have remained reasonably stable. Using rougher categories also helps relax the assumption that each level can only be achieved through one route through the education system, as individuals are now allowed to 'move freely' within the rough categories. The stylized system is presented in Figure 6.2. The simplified representation of the Dutch education system assumes that all respondents complete primary education (LO). After this, they face a choice between leaving the schooling system and continuing[4]. If they opt for the latter choice, they have to choose between the 'high track' (HAVO/VWO, that is, senior general secondary education and pre-university education) and the 'low track' (LBO/MAVO, that is, junior vocational education and junior general secondary education). Once they have finished their second diploma in either track they can choose whether or not to get a third diploma, continuing with: MBO (senior secondary vocational education) if they are in the low track, or HBO/WO (higher professional education and university) if they are in the high track.

[4]Since I measure education as the highest finished level of education, continuing education actually means continuing *and* finishing a subsequent level of eduction. Even though continuing education after primary education was compulsory during almost the entire historical period that is being studied, finishing a subsequent level of education was not compulsory.

Figure 6.3: Cohorts covered by each survey (survey numbers refer to the data references)

6.4.2 The data

The data were obtained from the International Stratification and Mobility File (ISMF) (Ganzeboom and Treiman, 2009). The ISMF now contains 55 surveys on the Netherland, carried out between 1958 and 2006. These were merged to increase the time period covered and the number of respondents, and to lessen the effect of individual surveys' idiosyncrasies. The cohorts covered by each survey are represented in Figure 6.3. It shows that information on the earliest and most recent cohorts primarily originates from a few surveys, while information on the middle cohorts originates from many surveys.

The purpose of this analysis is to compare the effect of a family's SES on the highest achieved level of education and on probabilities of passing transitions, both between men and women and across cohorts. Time was measured by the year in which the respondent was 12, scaled in decades since 1910. Information was available for the cohorts born between 1905 and 1991. Cohort is allowed to have a non-linear effect by representing it as a restricted cubic spline (Harrell, 2001; Royston and Parmar, 2002) as implemented in Stata (StataCorp, 2007) as the `mkspline` command. A restricted cubic spline means that the variable is split up at a minimum of three points

Figure 6.4: Number of observations per cohort

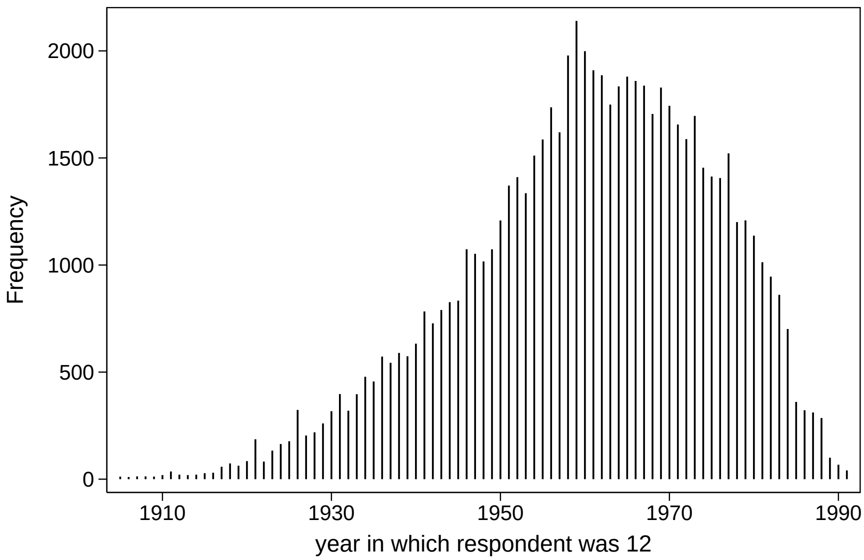

(or knots); in this case, cohort is split up at: 1920, 1950 and 1980. Between the first and the last knot the trend is represented by a cubic spline and before the first and after the last knot the trend is restricted to be linear. This restriction leads to a relatively stable non-linear curve. A family's SES was measured according to the father's score on the International Socio-Economic Index (ISEI) of occupational status (Ganzeboom and Treiman, 2003), as this measure was available for the largest number of cohorts. The original ISEI score is a continuous variable ranging from 10 to 90, but it was standardized to have a mean of 0 and a standard deviation of 1 for the cohort born in 1960 (approximately the cohort with the most observations in this study). Survey weights were used where available. The weighted number of respondents was 82,384, and after removing respondents with missing observations on any of the variables, 71,141 respondents remained.[5] The number of respondents was unequally distributed over the cohorts, as is shown in Figure 6.4.

A scale for the level of education was needed in order to estimate the relationship between the IEOpps and IEOut using equation (6.3). The scale that will be used in this example is similar to the one estimated in Chapter 3, which is estimated in such a

[5] Various Multiple Imputation models (Little and Rubin, 2002) were tried in Chapter 4 of this dissertation and none of them caused the conclusions to be changed.

way that it maximized the direct effect of education on income while controlling for the father's occupational status. This scale does not change over time, as I established in that chapter that even though the effect of education on occupational status changed over time, the scale of education remained constant. However, if evidence was found that the scale of education also changed over time, then such a changing scale could have easily been incorporated in the decomposition. For interpretability, the scale was coded in such a way that the mean was 0 and the variance was 1 for the cohort born in 1960.

6.4.3 Generalizing the decomposition to a tracked system

The model for the Dutch educational system as represented by Figure 6.2 is more complicated than the model in Figure 6.1, which was used to illustrate the decomposition of IEOut into IEOpps and weights. Whereas the model used in the example consists of a sequence of decisions to either continue or to stop, the model for the Dutch system also contains a 'branching point', or a choice between tracks. In this sense the model is akin to those proposed by Lucas (2001) and Breen and Jonsson (2000). This raises the question of whether the decomposition still holds in the more complicated model. For that reason the decomposition is derived again for the more complicated model. As before, logistic regressions were used to model the probabilities of passing the different transitions. Again, the IEOpp and the predicted probabilities belonging to transition k are represented by λ_k and $\hat{p}_{ki}$ respectively. The predicted level of education is now represented by equation (6.4).

$$
\begin{aligned}
E(L_i) \quad = \quad & (1-\hat{p}_{1i})l_1 + \\
& \hat{p}_{1i}(1-\hat{p}_{2i})(1-\hat{p}_{3i})l_2 + \\
& \hat{p}_{1i}(1-\hat{p}_{2i})\hat{p}_{3i}l_3 + \\
& p_{1i}\hat{p}_{2i}(1-\hat{p}_{4i})l_4 + \\
& p_{1i}\hat{p}_{2i}\hat{p}_{4i}l_5
\end{aligned}
\tag{6.4}
$$

Recall that the IEOut is first derivative of equation (6.4) with respect to a family's SES. This derivative is shown in equation (6.5).

$$
\begin{array}{llllllll}
\frac{\partial E(L_i)}{\partial SES} & = & & & & & & \\
\{1 & \times & \hat{p}_{1i}(1-\hat{p}_{1i}) & \times & [(1-\hat{p}_{2i})(1-\hat{p}_{3i})l_2+ & & & \\
 & & & & (1-\hat{p}_{2i})\hat{p}_{3i}l_3+ & & & \\
 & & & & \hat{p}_{2i}(1-\hat{p}_{4i})l_4+ & & & \\
 & & & & \hat{p}_{2i}\hat{p}_{4i}l_5 - l_1] & \} & \lambda_1 & + \\
\{\hat{p}_{1i} & \times & \hat{p}_{2i}(1-\hat{p}_{2i}) & \times & [(1-\hat{p}_{4i})l_4+\hat{p}_{4i}l_5- & & & \\
 & & & & (1-\hat{p}_{3i})l_2-\hat{p}_{3i}l_3] & \} & \lambda_2 & + \\
\{\hat{p}_{1i}(1-\hat{p}_{2i}) & \times & \hat{p}_{3i}(1-\hat{p}_{3i}) & \times & [(l_3-l_2)] & \} & \lambda_3 & + \\
\{\hat{p}_{1i}\hat{p}_{2i} & \times & \hat{p}_{4i}(1-\hat{p}_{4i}) & \times & [(l_5-l_4)] & \} & \lambda_4 &
\end{array}
\tag{6.5}
$$

Just as with the example described in section 6.3, IEOut is a weighted sum of the IEOpps, the λ_ks. The weights (the parts between curly brackets) consist of the same three parts:

1. The proportion of people at risk (1, $\hat{p}_{1i}$, $\hat{p}_{1i}(1-\hat{p}_{2i})$, and $\hat{p}_{1i}\hat{p}_{2i}$ respecively).

2. A part ($\hat{p}_{ki}(1-\hat{p}_{ki})$) that is small if virtually everybody passes or fails that transition and is largest when the probability of passing is 0.5.

3. The differences between the expected levels of education of those who pass the transitions and those who do not (these are the parts in the square brackets).

This case illustrates that the relationship between IEOut and IEOpp can be extended to tracked education systems. Using the same logic, the result can be extended to even more complex systems, such as those with more than two tracks. In that case a multinomial logit would be used to estimate the IEOpp. The Stata (StataCorp, 2007) package `seqlogit` (Buis, 2007b), which implements the decomposition, applies to this general version of the sequential logit model. The only limitation is that if one uses data with only the highest achieved level of education, one must ensure that for these more complicated systems, each level can only be reached through one — and only one — path through the education system.

6.4.4 Results

The following analysis consists of three parts. First, a descriptive analysis is performed on the differences in transition probabilities between men and women, and between cohorts. Second, the sequential response model described in the previous section is estimated. The results from this model are used to compute the IEOpps, the weights and the IEOut. Together these provide a detailed picture of status educational inequality and how it is influenced by educational expansion and gender inequality.

Third, the relationship between the transition probabilities and the weights is investigated in more detail by looking at the three components of the weights: the proportion at risk, the closeness of the transition probability to 50%, and the expected increase in the level of education when passing a transition.

The distribution of the highest achieved level of education is shown in Figure 6.5, for both males and females and for different cohorts. The changes over cohorts were smoothed using the `proprcspline` package (Buis, 2009a) in Stata (StataCorp, 2007). As with most other countries, the Netherlands experienced a period of educational expansion during the twentieth century. The proportion of pupils who only achieved LO (primary education) dropped dramatically, while the proportion attaining HBO/WO (higher professional and university) education and MBO (higher secondary vocational) strongly increased. Figure 6.5 also shows that MBO is a recent level of education. Whereas no one from the earlier cohorts completed this level of education, MBO completion has rapidly grown to about 40%. Furthermore, women experienced all of these developments later than men.

Figure 6.5: Distribution of highest achieved level of education for men and women over cohorts

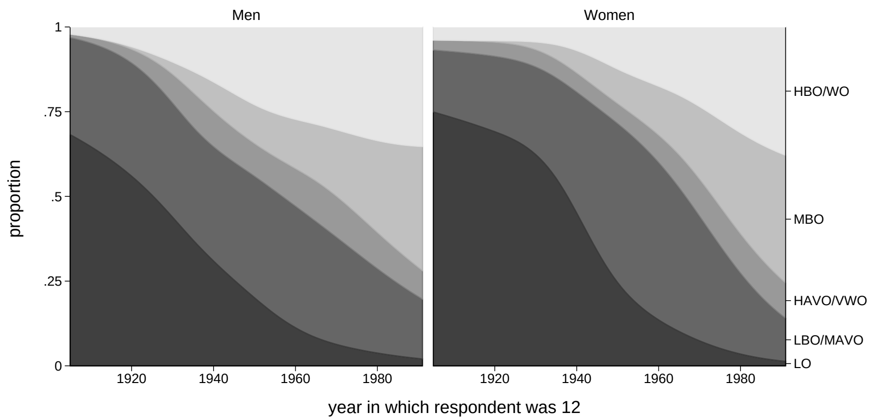

To investigate the IEOpps and IEOut and how they are influenced by gender and educational expansion (differences in the distribution of education between men and women and between cohorts respectively), sequential logit models were estimated separately for both men and women. The other variables are: cohort measured as a restricted cubic spline with knots at 1920, 1950, and 1980; the father's occupational status; and an interaction term with cohort. A model with a non-linear interaction between the father's occupational status and cohort was also estimated using the same restricted cubic spline as the main effect of cohort, but the non-linear terms proved to be non-significant (χ^2=4.73 with 4 df for men and χ^2=5.50 with 4 df for women). The results of this model are shown in Tables 6.1 and 6.2. The effects are log-odds ratios. The main effects of the father's occupational status are the IEOpps for the cohort born in 1910. This shows that the IEOpps for the higher transitions (in particular LBO/MAVO versus MBO and HAVO/VWO versus HBO/WO) are smaller than for the the lower transitions. This pattern has also been found by many other studies using sequential response models (Mare, 1980; Shavit and Blossfeld, 1993). Two explanations are commonly given for this phenomenon. First, persons passing the higher transitions are on average older than persons passing the lower transitions, and older persons are less likely to be influenced by their parents than younger persons (Shavit and Blossfeld, 1993). Second, selection on unobserved variables is likely to induce a negative correlation between the observed and unobserved variables, thus suppressing the effect of the observed variables at the higher transitions (Mare, 1981) (although Cameron and Heckman (1998) show that this does not always have to be the case). The interaction terms represent the change in effect for every ten-year change in cohort. These show that the effect of the father's occupational status changed most for the first transition. For men, this is the only transition in which the IEOpp changed significantly over cohorts. This pattern has already been found in the Netherlands (De Graaf and Ganzeboom, 1993), and is being found more frequently in studies of other countries (Breen and Jonsson, 2005).

From these results, one can derive predicted levels of education for each level of the father's occupational status, forming a non-linear regression line. Figure 6.6 presents these lines for three cohorts (1910, 1950, and 1990), and for men and women. The slope of this regression line will reveal how much the expected level of education changes when the father's occupational status changes by one unit, thus providing the IEOut. This slope is evaluated at the average father's occupational status. The father's occupational status is standardized, so a respondent with a typical background has a father's status of 0[6]. This figure shows that in all cases, having a father with a higher socioeconomic status will lead to a higher expected level of education. Also, it shows that while women initially suffered a disadvantage, they have overtaken men in the most recent cohort. Finally, the results show that for the earliest cohort, the inequality of educational outcomes for a respondent with a typical background was relatively

[6]However, the standardization uses the cohort born in 1960, and the average of the father's status increased over cohorts. The average of father's occupational status remained reasonably constant until about 1930 at about -0.2 and then steadily increased to 0.5. These changes not only reflect changes in economic structure, but also changes in the difference in the number of respondents between higher and lower status fathers. Consequently, it is hard to give a substantive interpretation to these changes. To simplify the analysis, a respondent with a typical background will be fixed at the typical background (average father's occupational status) for a typical cohort (1960).

Table 6.1: Sequential response model for men

	LO v more	LBO/MAVO v HAVO/VWO	LBO/MAVO v MBO	HAVO/VWO v HBO/WO
Father's status	0.912	0.694	0.263	0.446
	(15.28)	(14.19)	(3.44)	(5.91)
Father's status X of Cohort	-0.068	-0.015	-0.004	-0.033
	(-5.09)	(-1.62)	(-0.30)	(-2.35)
RC spline term 1 of Cohort	0.566	0.316	0.461	0.461
	(17.54)	(9.15)	(9.45)	(7.93)
RC spline term 2 of Cohort	-0.000	0.013	0.002	0.015
	(-0.01)	(7.08)	(0.97)	(4.82)
Constant	-0.590	-1.470	-2.893	-0.806
	(-6.36)	(-13.13)	(-18.00)	(-4.24)
N	43770			
Log likelihood	-50032.082			

z statistics in parentheses

Table 6.2: Sequential response model for women

	LO v more	LBO/MAVO v HAVO/VWO	LBO/MAVO v MBO	HAVO/VWO v HBO/WO
Father's status	0.874	1.021	0.412	0.079
	(15.33)	(17.23)	(5.21)	(0.88)
Father's status X cohort	-0.068	-0.063	-0.021	0.029
	(-5.34)	(-6.00)	(-1.51)	(1.82)
RC spline term 1 of Cohort	0.743	0.103	0.129	0.345
	(21.26)	(2.27)	(2.33)	(4.67)
RC spline term 2 of Cohort	-0.001	-0.008	-0.022	0.008
	(-0.24)	(-3.58)	(-8.27)	(2.36)
Constant	-1.727	-1.693	-2.431	-0.760
	(-17.05)	(-10.88)	(-12.87)	(-2.99)
N	43675			
Log likelihood	-45830.33			

z statistics in parentheses

small (the curve is rather flat), because everybody in the immediate neighbourhood of the respondent with an average family background had an expected level of education that was close to the minimum. However, in this same cohort, respondents with very high-status parents do a lot better than the other respondents, which would lead to a high inequality of educational outcome. In other words, in this chapter estimates of the local educational inequality will be obtained, and if one were to estimate a measure of global educational inequality instead, the estimate would be higher for the earliest cohorts.

Figure 6.6: Expected highest achieved level of education according to the sequential logit model

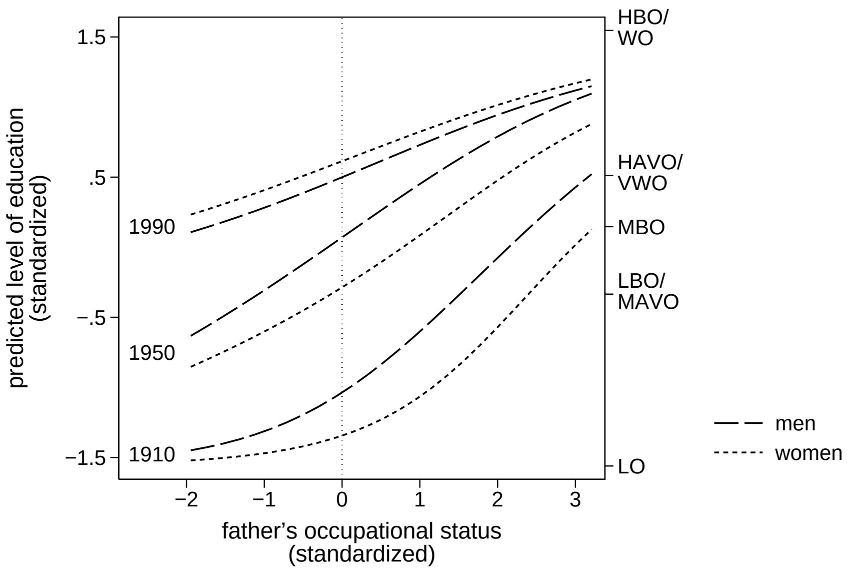

Figure 6.7 shows the estimates of IEOut that have been derived from the sequential logit model. Both education and the father's occupational status are scaled in such a way that the mean for the cohort 1960 is 0 and the standard deviation is 1. So this measure of IEOut is similar to a standardized regression coefficient. IEOut displays two striking features: the first is the trend in IEOut, which initially increases and then decreases. The second feature is the initially lower IEOut for women. These are not unique to the sequential logit model, since in Chapter 4 I found similar patterns using different methods. In order to explain these patterns IEOut will be broken down into its components, in three steps.

Figure 6.7: IEOut according to the sequential logit model

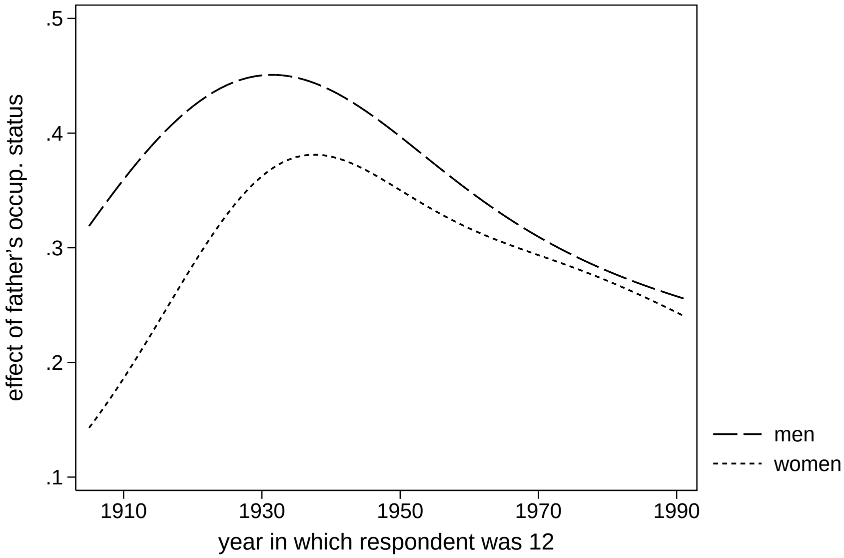

The first step looks at the contributions of each transition to IEOut. The IEOut is a weighted sum of each transition's IEOpp, so each transition contributes the amount of weight times IEOpp to IEOut. This is shown in Figure 6.8. A striking feature is that the final two transitions (HAVO/VWO to HBO/WO and LBO/MAVO to MBO) contribute negligible amounts to IEOut. Furthermore, the initial increase and later decrease in IEOut seems to be primarily the result of what happened at the first transition. Finally, there has been a shift between the first and the second transitions as the dominant source of IEOut.

The second step consists of breaking up each transition's contribution into its two parts: the weight and the IEOpp. Since the contribution is the product of these two terms, it can be visualized as the area of a rectangle, with a height equal to the IEOpp and a width equal to the weight. For men and women, this is shown in Figures 6.9 and 6.10. The horizontal axis shows the weights and the vertical axis the IEOpp, while the columns represent the cohorts and the rows represent the transitions. These figures show that the initial increase in the contribution of the first transition is due to an increase in its weight, while the later decrease of this transition is due to both a decrease in the weight and a decrease in the IEOpp. The increase in importance of the second transition is entirely due to the increase in the weight of this transition.

Figure 6.8: Contribution of each transition to IEOut

For women, this increase in weight actually offsets a decrease in IEOpp. The low contributions of both higher transitions are due to both low IEOpp and low weight.

The third step breaks the weights down into their three components. Figure 6.11 (a) shows the changes in the weights over time in more detail. The changes in these weights capture the consequences of changes in the distribution of education on IEOut. These weights are the product of three components: the proportion of people at risk at each transition (Figure 6.11 (b)); the closeness to 50% of the proportion of people passing (the variance) (Figure 6.11 (c)); and the difference in the expected level of education between those passing and those failing a transition (Figure 6.11 (d)). Figure 6.11 shows that the initial increase and the later decline in the first transition's influence is primarily due to the variance. Initially, any inequality at the first transition affected few people, because a low proportion passed. As the proportion of people passing increased, the transition received more weight, until half of the students passed, after which inequality affected less people again because few people failed. The increase in importance of the second transition is partly due to the variance, but also to a strong increase in the number of students that are at risk of making this transition. Notice that these developments at the first two transitions provide a substantively interpretable mechanism through which educational expansion influences IEOut. For women, these developments have occurred later, leading initially to smaller weights. The last two transitions receive relatively small weights because relatively few people are at risk of passing these transitions, and those who pass gain relatively little. Those who pass the first two transitions gain both the immediate increase in level of education and the possibility of gaining an extra level of education (either MBO or

HBO/WO), while in the third and fourth transition, people gain only the immediate increase in level of education.

Figure 6.9: Decomposition of IEOut into IEOpps and weights

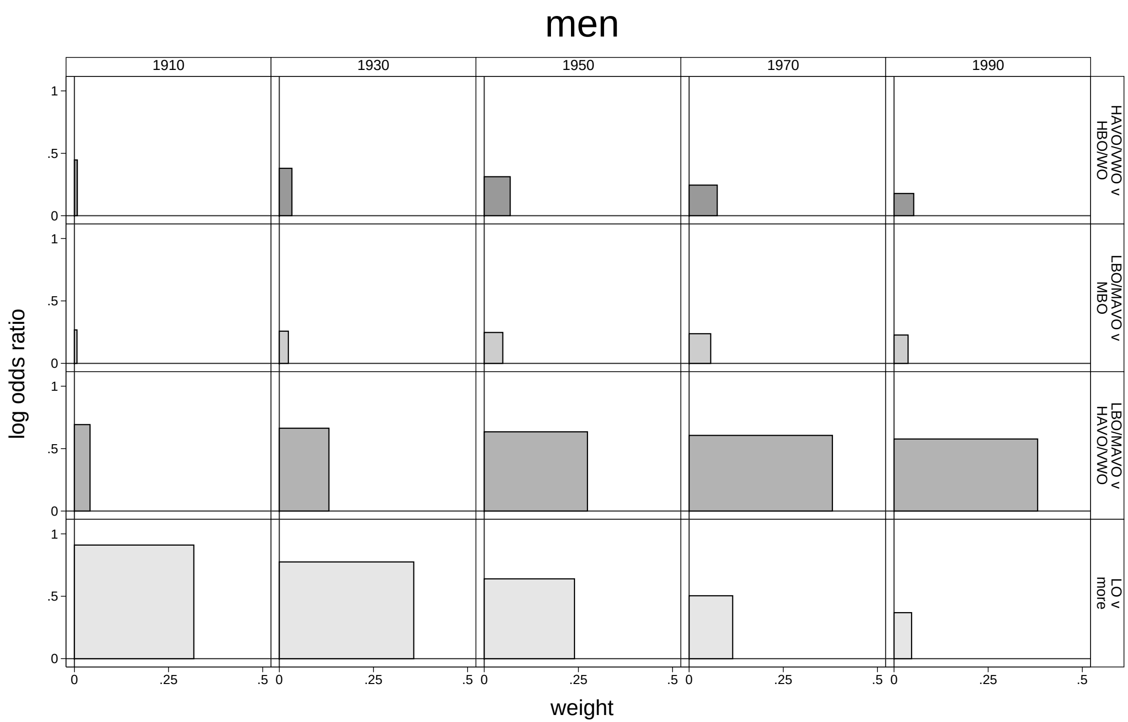

Figure 6.10: Decomposition of IEOut into IEOpps and weights

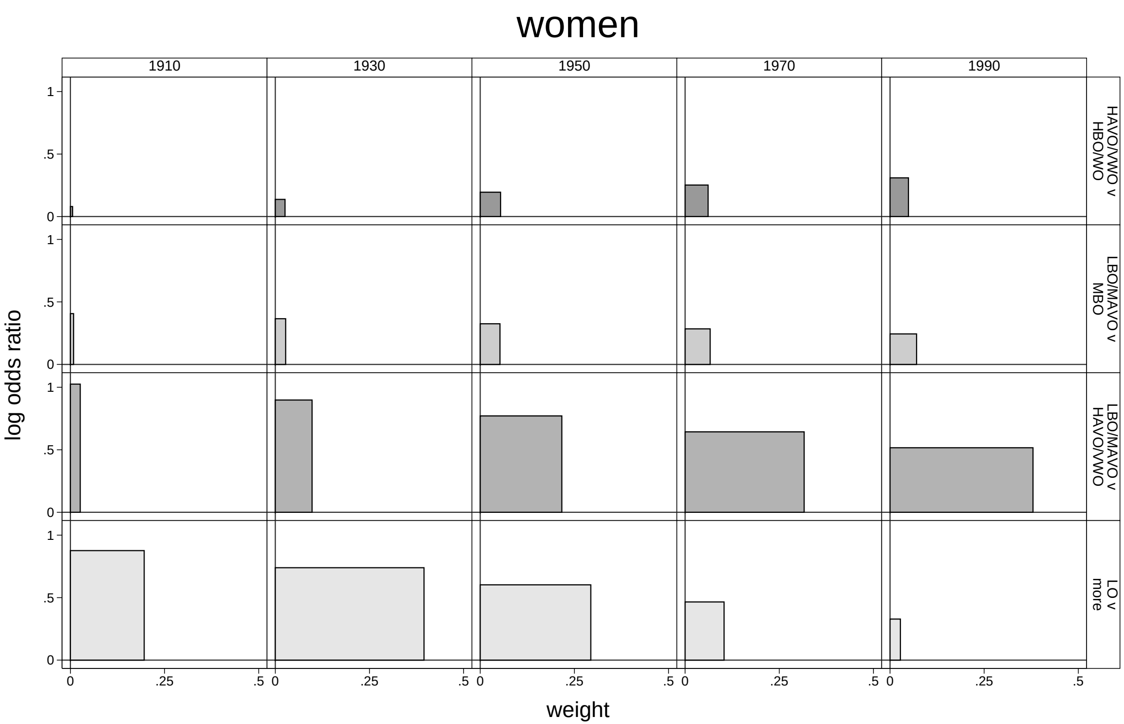

Figure 6.11: Weights and their components

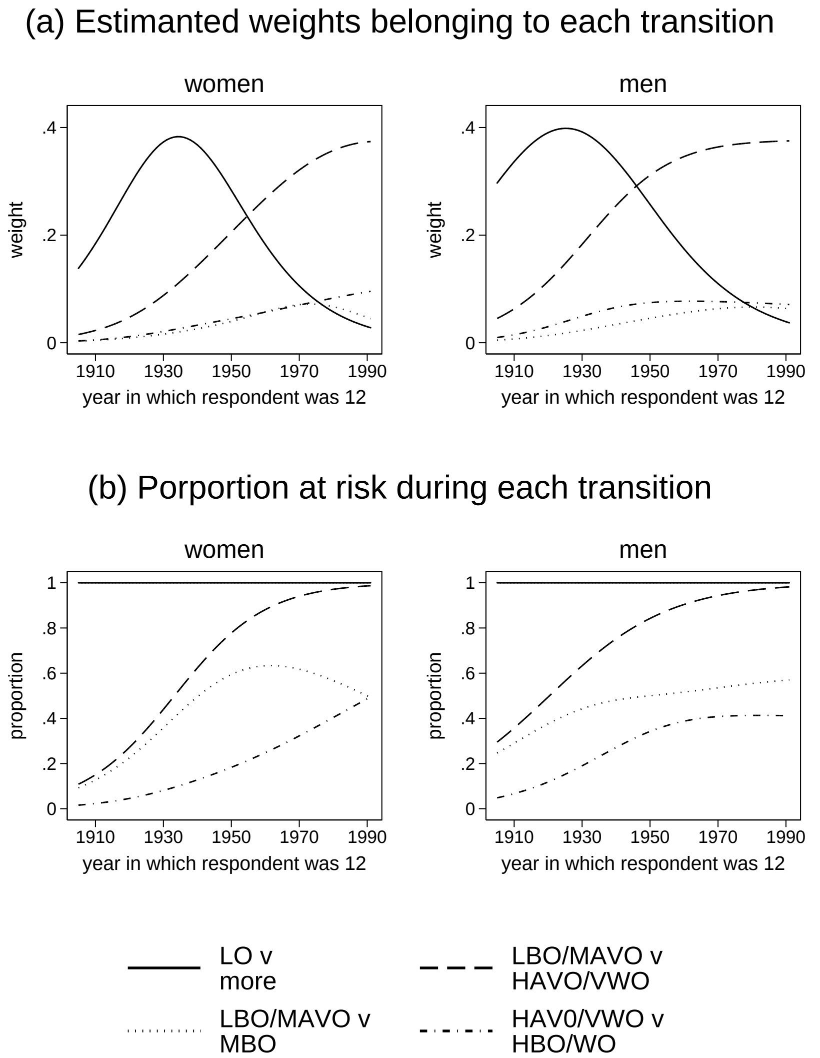

Figure 6.11: Weights and their components (continued)

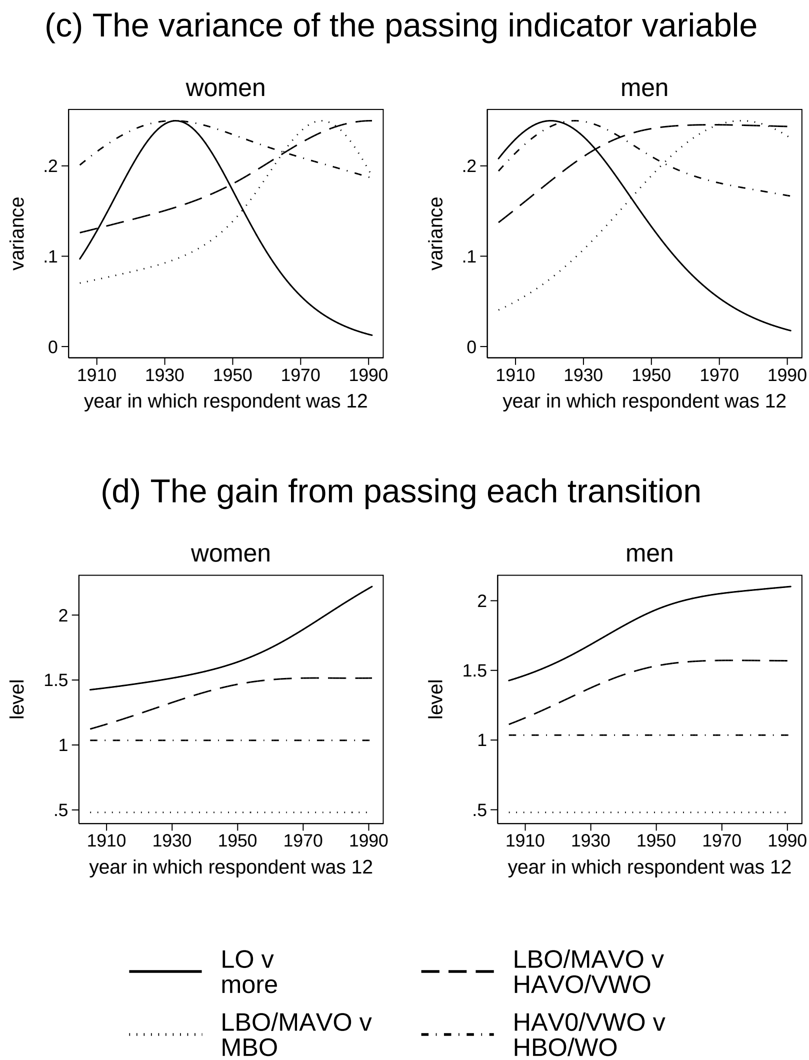

6.5 Conclusion

This chapter began by making a distinction between two types of inequality of educational opportunity (IEO): inequality of educational opportunities during the process of attaining education, which I called Inequality of Educational Opportunities proper (IEOpp), and inequality of educational opportunities in terms of the outcome of the educational process, which I called Inequality of Educational Outcomes (IEOut). Mare (1981) demonstrated that differences in IEOut across cohorts (or other groups) depend on both the differences in IEOpp and differences in the distribution of education. However, this literature did not study the relationship between IEOpp, IEOut and the distribution of education, but instead treated this relationship as a 'black box'. This was used as an argument for studying only IEOpps and for controlling for the distribution of education rather than of studying its effects. This chapter seeks to change this by answering the following two questions:

- How are IEOut and IEOpp related to one another, and how can this relation be used for a meaningfully integrated analysis of IEOpp and IEOut?
- How are IEOut and the distribution of education related to one another, and how can this relation be used for an analysis of the influence of changes in the distribution of education on IEOut?

The first question is based on the observation that IEOpp and IEOut are not competing descriptions of IEO but natural complements, because a description of a process (the IEOpps) and a description of the outcome of that process (the IEOut) are natural complements. Treating IEOpps and IEOut as complementary creates the challenge to move beyond a separate discussion of these two estimates to an integrated discussion of IEOpp and IEOut. The second question is based on the observation that the influence of changes in the distribution of education on estimates of IEO is a phenomenon of substantive interest. One such change in the distribution of education is the general increase in highest achieved level of education over cohorts, which is one of the most universal and far-reaching changes in educational systems across countries during the 20^{th} century (Hout and DiPrete, 2006). The consequences for IEO of such a major change in the educational system deserve to be studied rather than just controlled for.

These questions are answered by showing that the sequential logit model, which was proposed by Mare (1981) for estimating IEOpps, also implies an estimate for IEOut. This estimate of IEOut is a weighted sum of IEOpps such that an IEOpp that belongs to a certain transition between levels of education receives more weight if more people are at risk of passing that transition; if passing or failing the transition is less universal (that is, if the proportion of respondents who pass is closer to 50%); and if there is a larger difference in the expected level of education between people who pass and fail that transition. This decomposition shows how IEOpp and IEOut are related and allows for an integrated discussion of these two by showing to what extent each transition's IEOpp contribute to IEOut. The weights also allows one to study the impact of changes in the distribution of education on IEOut, as these weights depend on the distribution in a substantively interpretable way.

The application of this decomposition was illustrated using an analysis of changes in IEO in the Netherlands between 1905 and 1991. It showed that the composition of IEOut shifted from being primarily determined by the IEOpp of the first transition (whether or not to continue after primary education) to being primarily determined by the IEOpp of the second transition (the choice between the vocational and the academic track). The IEOpps of the later transitions contributed relatively little to IEOut throughout the period being studied. The differences in the distribution of education across cohorts (educational expansion) and gender (gender educational inequality) were shown to explain this shift in importance between the first and second transitions and two main features of the trend in IEOut. First, the trend over cohorts showed an initial increase followed by a decrease. Second, the IEOut is initially lower for women. The initial increase in IEOut can be explained by the increase in the proportion of students that pass the first two transitions from less than 50% to around 50%, thus initially increasing the weights for both transitions. The weight for the second transition also increased as more students became at risk of passing that transition. The subsequent decrease in IEOut happened because the weight of the first transition's IEOpp sharply decreased since passing that transition became near universal. These changes also explain the shift in importance between the IEOpps of the first and second transitions. The decrease in the difference between men and women in IEOut was caused by the fact that initially fewer women passed each transition, causing each transition's weight to be less for women than for men. For the later cohorts, weights were approximately equal between men and women, because women were as likely as men —or even more likely — to pass transitions, thus causing a convergence in IEOut of men and women.

This chapter defined IEOut in such a way that it is meaningfully influenced by changes in the distribution of education. There is however an important body of research in this literature that uses log-linear models that summarize the IEOut in a single odds ratio (De Graaf and Ganzeboom, 1990; Ganzeboom and Luijkx, 2004a,b). Unlike the measure of IEOut used in this chapter, the odds ratio controls changes in the distribution of education, that is, educational expansion. I would argue that this is not necessarily a good thing: changes in IEOut over time are studied not because we think that time directly influences IEOut, but that society changes over time and these changes lead to changes in IEOut. The aim of such an analysis should be to study how these changes in society influenced IEOut, not sweep them under the carpet by controlling for them.

In future research, the decomposition presented in this chapter can be generalized in a number of ways. First, the decomposition can be applied to some models that have been proposed to address the critique on the sequential logit model by Cameron and Heckman (1998). The decomposition can be applied to those models that are direct adaptations of the sequential logit model (for example: Mare 1993, 1994; and Chapter 7 of this dissertation), but not to models that do not use the (multinomial) logit link function (for example Lucas et al., 2007; Holm and Jæger, 2008). Second, the decomposition requires that each level of education is assigned a value. In this chapter, these values are constant over time, but there has been debate on whether the values of educational categories have changed as a consequence of strong changes in the distribution of education and the labor market (Rumberger, 1981; Clogg and Shockey, 1984; Groot and Maassen van den Brink, 2000). If one has time-varying estimates of the value of the levels of education, then these could also be incorporated in the decomposition. Changes in these values would influence IEOut through only one of the three components of the weight: the difference in the expected highest achieved level of education between people who pass and fail a transition. The decomposition could thus also be used to study the impact of possible changes in the values of educational levels. Third, the analysis is based on data on the highest achieved level of education in combination with a stylized model of the education system. The transitions that respondents have passed were derived from these two pieces of information rather than being directly observed. The main advantage of using highest achieved levels of education is that much more data is available on the highest achieved level of education and that this data covers a larger period than data on actual transitions. However, an additional analysis using observed transitions is desirable. An interesting question that could be answered this way would be the impact of 'second chance paths', that is, paths where one switches from one track to another. The effect of these second chance paths on IEO is not clear: on the one hand these second chance paths could offer a way out of lower tracks for those disadvantaged students that were disproportionably assigned to them, on the other hand students from advantaged background are generally better capable of making the best use of these 'loopholes'. An additional advantage of using observed transitions is that one no longer has to rely on pseudo-cohorts to measure trends over time, as in that case one directly observes when a transition occurred.

In conclusion, this chapter has shown how the study of educational inequality can be enriched by studying IEOpp and IEOut as complementary pieces of information and by studying the impact of the distribution of education, rather than by simply controlling for it. This has the key advantage of enabling an integrated discussion of IEOpp and IEOut and the study of the impact of phenomena such as educational expansion.

Appendix: Derivation of equation (6.3)

Equation (6.3) is the first derivative of equation (6.2). Equation (6.2) is repeated below:

$$E(L_i) = (1 - \hat{p}_{1i})l_0 + \hat{p}_{1i}(1 - \hat{p}_{2i})l_1 + \hat{p}_{1i}\hat{p}_{2i}(1 - \hat{p}_{3i})l_2 + \hat{p}_{1i}\hat{p}_{2i}\hat{p}_{3i}l_3$$

whereby the $\hat{p}_{ki}$s are represented by equation (6.1), repeated below:

$$\hat{p}_{ki} = \frac{\exp(\alpha_k + \lambda_k SES_i)}{1 + \exp(\alpha_k + \lambda_k SES_i)} \quad \text{if} \quad y_{k-1\,i} = 1$$

This derivative can be computed using the sum rule,[7] the product rule,[8], and the derivative of a logistic regression equation.[9] Using the sum rule, the first derivative can be written as:

[7]Suppose that we have two functions of SES: $f(SES)$ and $g(SES)$. The sum rule states that the derivative of the sum of these functions with respect to SES is (e.g. Gill, 2006, p. 190):

$$\frac{\partial(f(SES) + g(SES))}{\partial SES} = \frac{\partial f(SES)}{\partial SES} + \frac{\partial g(SES)}{\partial SES}$$

[8]The product rule states that the derivative of the product of these functions with respect to SES is (e.g. Gill, 2006, p. 191):

$$\frac{\partial(f(SES) \times g(SES))}{\partial SES} = \frac{\partial f(SES)}{\partial SES} g(SES) + \frac{\partial g(SES)}{\partial SES} f(SES)$$

A special case occurs when a function of SES is multiplied by a constant c because the first derivative of a constant is zero:

$$\frac{\partial(cf(SES))}{\partial SES} = \frac{\partial f(SES)}{\partial SES} c + \frac{\partial c}{\partial SES} f(SES) = \frac{\partial f(SES)}{\partial SES} c$$

[9]Equation (6.1) is a logistic regression equation, which has a known first derivative (e.g. equation 3.14 Long, 1997):

$$\frac{\partial \hat{p}_{ki}}{\partial SES} = \hat{p}_{ki}(1 - \hat{p}_{ki})\lambda_k$$

Together with the sum and the product rule this also implies that:

$$\begin{aligned}
\frac{\partial(1 - \hat{p}_{ki})}{\partial SES} &= \frac{\partial 1}{\partial SES} + \frac{\partial - \hat{p}_{ki}}{SES} \quad \text{(sum rule)} \\
&= -\frac{\partial \hat{p}_{ki}}{SES} \quad \text{(product rule)} \\
&= -\hat{p}_{ki}(1 - \hat{p}_{ki})\lambda_k
\end{aligned}$$

$$\frac{\partial E(L_i)}{\partial SES} = \frac{\partial(1-\hat{p}_{1i})l_0}{\partial SES} + \frac{\partial p_1(1-\hat{p}_{2i})l_1}{\partial SES} + \frac{\partial \hat{p}_{1i}\hat{p}_{2i}(1-\hat{p}_{3i})l_2}{\partial SES} + \frac{\partial \hat{p}_{1i}\hat{p}_{2i}\hat{p}_{3i}l_3}{\partial SES}$$

Using the product rule, this can be rewritten as:

$$\begin{aligned}
\frac{\partial E(L_i)}{\partial SES} \quad = \quad & l_0 \frac{\partial(1-\hat{p}_{1i})}{\partial SES} + \\
& l_1 \left(\frac{\partial \hat{p}_{1i}}{\partial SES}(1-\hat{p}_{2i}) + \frac{\partial(1-\hat{p}_{2i})}{\partial SES}\hat{p}_{1i} \right) + \\
& l_2 \left(\frac{\partial \hat{p}_{1i}}{\partial SES}\hat{p}_{2i}(1-\hat{p}_{3i}) + \frac{\partial \hat{p}_{2i}}{\partial SES}\hat{p}_{1i}(1-\hat{p}_{3i}) + \frac{\partial(1-\hat{p}_{3i})}{\partial SES}\hat{p}_{1i}\hat{p}_{2i} \right) + \\
& l_3 \left(\frac{\partial \hat{p}_{1i}}{\partial SES}\hat{p}_{2i}\hat{p}_{3i} + \frac{\partial \hat{p}_{2i}}{\partial SES}\hat{p}_{1i}\hat{p}_{3i} + \frac{\partial \hat{p}_{3i}}{\partial SES}\hat{p}_{1i}\hat{p}_{2i} \right)
\end{aligned}$$

All derivatives in the equation are derivatives of logistic regression equations. To facilitate the comparison with the previous equation, curly brackets are used to enclose these derivatives.

$$\begin{array}{ll}
\frac{\partial E(L_i)}{\partial SES} & = \\
l_0 & \{-\hat{p}_{1i}(1-\hat{p}_{1i})\lambda_1\}+ \\
l_1 & (\{\hat{p}_{1i}(1-\hat{p}_{1i})\lambda_1\}(1-\hat{p}_{2i}) + \{-\hat{p}_{2i}(1-\hat{p}_{2i})\lambda_2\}\hat{p}_{1i}) + \\
l_2 & (\{\hat{p}_{1i}(1-\hat{p}_{1i})\lambda_1\}\hat{p}_{2i}(1-\hat{p}_{3i}) + \{\hat{p}_{2i}(1-\hat{p}_{2i})\lambda_2\}\hat{p}_{1i}(1-\hat{p}_{3i})+ \\
 & \{-\hat{p}_{3i}(1-\hat{p}_{3i})\lambda_3\}\hat{p}_{1i}\hat{p}_{2i}) + \\
l_3 & (\{\hat{p}_{1i}(1-\hat{p}_{1i})\lambda_1\}\hat{p}_{2i}\hat{p}_{3i} + \{\hat{p}_{2i}(1-\hat{p}_{2i})\lambda_2\}\hat{p}_{1i}\hat{p}_{3i}+ \\
 & \{\hat{p}_{3i}(1-\hat{p}_{3i})\lambda_3\}\hat{p}_{1i}\hat{p}_{2i})
\end{array}$$

The terms in this equation can be rearranged in such a way that all elements that have the same IEOpp (λ_k) in common are grouped together.

$$\begin{array}{ll}
\frac{\partial E(L_i)}{\partial SES} & = \\
\lambda_1 & \{-\hat{p}_{1i}(1-\hat{p}_{1i})l_0 + \hat{p}_{1i}(1-\hat{p}_{1i})(1-\hat{p}_{2i})l_1 + \\
 & \hat{p}_{1i}(1-\hat{p}_{1i})\hat{p}_{2i}(1-\hat{p}_{3i})l_2 + \hat{p}_{1i}(1-\hat{p}_{1i})\hat{p}_{2i}\hat{p}_{3i}l_3\}+ \\
\lambda_2 & \{-\hat{p}_{2i}(1-\hat{p}_{2i})\hat{p}_1 l_1 + \hat{p}_{2i}(1-\hat{p}_{2i})\hat{p}_1(1-\hat{p}_{3i})l_2 + \\
 & \hat{p}_{2i}(1-\hat{p}_{2i})\hat{p}_{1i}\hat{p}_{3i}l_3\}+ \\
\lambda_3 & \{-\hat{p}_{3i}(1-\hat{p}_{3i})\hat{p}_{1i}\hat{p}_{2i}l_2 + \hat{p}_{3i}(1-\hat{p}_{3i})\hat{p}_{1i}\hat{p}_{2i}l_3\}
\end{array}$$

Simplifying this equation will yield equation (6.3).

Chapter 7

The consequences of unobserved heterogeneity in a sequential logit model

7.1 Introduction

Many processes can be described as a nested sequence of decisions or steps. Consider the three following examples. Mare (1979, 1980, 1981) describes the process of attaining education as the result of a sequence of transitions between educational levels, for example: 1) whether to finish secondary education or to leave school with only primary education, and 2) whether or not to finish tertiary education given that one finished secondary education. O'Rand and Henretta (1982) describe the decision when to retire using the following sequence of decisions: 1) whether to retire before age 62 or later, and 2) whether to retire before age 64 or later given that one has not retired before age 62. Cragg and Uhler (1970) describe the demand for automobiles as the result of the following sequence of decisions: 1) whether or not to buy an automobile, 2) whether to add an automobile or to replace an automobile given that one decided to buy an automobile, 3) whether or not to sell an automobile or not given that one decided not to buy an automobile. An attractive model for these processes is to estimate a separate logistic regression for each step or decision. These steps or decisions are often called transitions. This model is known under a variety of names: sequential response model (Maddala, 1983), sequential logit model (Tutz, 1991), continuation ratio logit (Agresti, 2002), model for nested dichotomies (Fox, 1997), and the Mare model (Shavit and Blossfeld, 1993). This model has however been subject to an influential critique by Cameron and Heckman (1998). Their main point starts with the observation that the sequential logit model, like any other model, is a simplification of reality and will not include all variables that influence the probability of passing a transition. The presence of these unobserved variables is often called unobserved heterogeneity, and it will lead to biased estimates, even if these unobserved variables are not confounding variables. There are two mechanisms through which these unobserved non-confounding variables will influence the results. The first mechanism, which I will call the averaging mechanism, is based on the fact that leaving a variable out of the model means that one models the probability of passing a transition averaged over the variable that was left out. The effect of the remaining variables on this average probability of passing a transition is not the same as the effect of these variables on the probability that an individual passes that transition, because the relationship between the variable left out of the model and the probability is non-linear (Neuhaus and Jewell, 1993; Cameron and Heckman, 1998; Allison, 1999). The second mechanism, which I will call the selection mechanism, is based on the fact that even if a variable is not a confounding variable at the initial transition because it is uncorrelated with any of the observed variables, it will become a confounding variable at the higher transitions because the respondents who are at risk of passing these higher transitions form a selected sub-sample of the original sample (Mare, 1980; Cameron and Heckman, 1998).

The aim of this chapter is to propose a sensitivity analysis with which one can investigate the consequences of unobserved non-confounding variables in a sequential logit model. This will be done by specifying a set of plausible scenarios concerning this unobserved variability and estimating the individual-level effects within each of these scenarios, thus creating a range of plausible values for the individual-level effects.

Any method for studying such individual-level effects will have to deal with the fact that it tries to control for variables that have not been observed. This is a problem that also occurs with other models that try to estimate causal effects (Holland, 1986). A common strategy in these causal models is to use information that might be available outside the data. The clearest example of this is the experiment in which one knows that the respondents have been randomly assigned to the treatment and the control group, and it is this information that is being used to control for any unobserved variables. Various variations on this strategy have been proposed for non-experimental settings (Morgan and Winship, 2007), for example one might know that a variable influences the main explanatory variable but not the outcome variable, in which case one can use this variable as an instrumental variable, or one might know that all variables influencing the main explanatory variable are present in the data, in which case one can use propensity score matching. An example of such a strategy that has been applied to the sequential logit model is the model by Mare (1993, 1994), who used the fact that siblings are likely to have a shared family background. If one has data on siblings, one can thus use this information for controlling for unobserved variables on the family level. Another example of this strategy is the model used by Holm and Jæger (2008), who use instrumental variables in a sequential probit model[1] to identify individual-level effects. The strength of this strategy depends on the strength of the information outside the data that is being used to identify the model. However, such external information is often not available. In those cases, one can still use these mod-

[1]The sequential probit model is similar to the sequential logit model except that the probit link function is used rather than the logit link function.

els, except that the identification is now solely based on untestable assumptions. This implies a subtle shift in the goal of the analysis: instead of trying to obtain an empirical estimate of a causal effect, one is now trying to predict what would happen if a certain scenario were true. This is not unreasonable: the causal effects are often the quantity of interest, and if it is not possible to estimate them, then the results of these scenarios are the next best thing. However, the modelling challenge now changes from making the best use of some information outside the data to finding the most informative comparison of scenarios. The goal of such an analysis is to find a plausible range of estimates of the causal effect and to assess how sensitive the conclusions are to changes in the assumptions (Rosenbaum and Rubin, 1983; Rosenbaum, 2002; DiPrete and Gangl, 2004). I will propose a set of scenarios that will allow one to directly manipulate the source of the problem: the degree of unobserved heterogeneity. This way one can compare how the results would change if there is a small, moderate, or large amount of unobserved heterogeneity.

This chapter will start with a more detailed discussion of how unobserved heterogeneity can cause bias in the estimates of the effect of the observed variables, even if the unobserved variables are initially non-confounding variables. I will then propose a sensitivity analysis, by specifying a series of scenarios concerning the unobserved variables. The estimation of the effects within these scenarios will be discussed next. Finally, the method will be illustrated by replicating an analysis of the effect of parental background on educational attainment in the Netherlands by De Graaf and Ganzeboom (1993) and in Chapter 2, and assessing how robust their results are for changes in assumptions about unobserved heterogeneity.

7.2 The sequential logit model and two effects of unobserved heterogeneity

The effect of unobserved heterogeneity in a sequential logit model is best explained using an example. Figure 7.1 shows a hypothetical process, which is to be described using a sequential logit model. There are three levels in this process: A, B and C. This process consists of two transitions: the first transition is a choice between A on the one hand and B and C on the other. The second transition is a choice between B and C for those who have chosen B and C in first transition. Figure 7.1 could be a representation of both the educational attainment example and the retirement example in the introduction. In the former case, A would correspond to primary education, B would correspond to secondary education, and C would correspond to tertiary education. In the latter case, A would correspond to retire before age 62, B would correspond to retire between age 62 and 64, and C would correspond to retire after age 64.

Figure 7.1: Hypothetical process

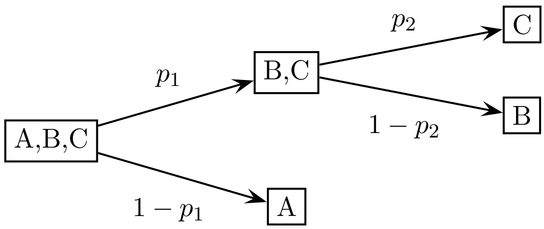

The sequential logit model models the probabilities of passing these transitions. This is done by estimating a logistic regression for each transition on the sub-sample that is at risk, as in equations (7.1) and (7.2). Equation (7.1) shows that the probability labelled p_1 in Figure 7.1 is related to two explanatory variables x and z through the function $\Lambda()$, while equation (7.2) shows the same for the probability labelled p_2 in Figure 7.1. The function $\Lambda()$ is defined such that $\Lambda(u) = \frac{\exp(u)}{1+\exp(u)}$. This function ensures that the predicted probability always remains between 0 and 1, by modelling the effects of the explanatory variables as S-shaped curves. The coefficients of x and z (β_{11}, β_{21}, β_{12}, and β_{22}) can be interpreted as log odds ratios, while the constants (β_{01} and β_{02}) represent the baseline log odds of passing the first and second transitions.

$$p_1 = \Pr(y \in \{B, C\}|x, z) = \Lambda(\beta_{01} + \beta_{11}x + \beta_{21}z) \quad (7.1)$$
$$p_2 = \Pr(y \in \{C\}|x, z, y \in \{B, C\}) = \Lambda(\beta_{02} + \beta_{12}x + \beta_{22}z) \quad (7.2)$$

Table 7.1 turns Figure 7.1 and equations (7.1) and (7.2) into a numerical example. Panel (a) shows the counts, the probabilities of passing, the odds and log odds ratios when z is observed, while panel (b) shows what happens in this example when z is not observed. Both x and z are dichotomous (where low is coded as 0 and high as 1), and during the first transition x and z are independent, meaning that z is not a confounding variable at the first transition. The sequential logit model underlying this example is presented in equations (7.3) and (7.4).

$$\Pr(y \in \{B, C\}|x, z) = \Lambda[\log(.333) + \log(3)x + \log(3)z] \quad (7.3)$$
$$\Pr(y \in \{C\}|x, z, y \in \{B, C\}) = \Lambda[\log(.333) + \log(3)x + \log(3)z] \quad (7.4)$$

Consider the first transition in panel (a). The constant in the logistic regression equation is the log odds of passing for the group with value 0 for all explanatory vari-

Table 7.1: Example illustrating the consequences of not observing a non-confounding variable (z)

(a) while observing z

transition	z	x	y: A	y: B, C	N	Pr(pass)	odds(pass)	log odds ratio
1	low	low	300	100	400	0.25	0.333	log(3)
		high	200	200	400	0.5	1	
	high	low	200	200	400	0.5	1	log(3)
		high	100	300	400	0.75	3	
			B	C				
2	low	low	75	25	100	0.25	0.333	log(3)
		high	100	100	200	0.5	1	
	high	low	100	100	200	0.5	1	log(3)
		high	75	225	300	0.75	3	

(b) without observing z

transition	x	y: A	y: B, C	N	Pr	odds	log odds ratio
1	low	500	300	800	0.375	0.6	log(2.778)
	high	300	500	800	0.625	1.667	
		B	C				
2	low	175	125	300	0.417	0.714	log(2.6)
	high	175	325	500	0.65	1.857	

ables, so the constant is in this case log(.333). The effect of x in a logistic regression equation is the log odds ratio. Within the low z group, the odds of passing for the low x group is .333 and the odds of passing for the high x group is 1, so odds ratio is $\frac{1}{.333} = 3$, and the log odds ratio is log(3). The effect of x in the high z group is also log(3), so there is no interaction effect between x and z. The effect of z can be calculated by comparing the odds of passing for a high z and a low z individual within the low x group, which results in a log odds ratio of log(3). There is no interaction between x and z, so the log odds ratio for z within the high x group is also log(3). Panel (b) shows what happens if one only observes x and y but not z. For example, in that case 300 + 200 = 500 low x persons are observed to have failed the first transition and 100 + 200 = 300 low x persons are observed to have passed the first transition. The resulting counts are used to calculate the probabilities, odds, and log odds ratios. Panel (b) shows that the log odds ratios of x are smaller than those computed in panel (a). Leaving z out of the model thus resulted in an underestimation of the effect of x for both the first and the second transition, even though z was initially uncorrelated with x.

This example can be used to illustrate both mechanisms through which unobserved heterogeneity can lead to biased estimates of the individual-level effects. First, the selection mechanism can explain part of the underestimation of the effect of x at the second transition. A characteristic of the sequential logit model is that even if z is not a confounding variable during the first transition, it will become a confounding variable during the later transitions (Mare, 1980; Cameron and Heckman, 1998). The example was created such that z and x are independent during the first transition, as the distribution of z is equal for both the low x group and high x group. As a consequence, z cannot be a confounding variable during the first transition. But this is no longer true during the second transition. For the high x group, the proportion of persons with a high z is $300/500 = .6$, while for the low x group that proportion is $200/300 = .667$. The selection at the first transition has thus introduced a negative correlation between x and z, and z has become a confounding variable. If one does not observe z, and thus can not control for z, one would expect to underestimate the effect of x at the second transition. This could in part explain the underestimation of the effect of x in the second transition in panel (b) of Table 7.1, but not the underestimation of the effect of x in the first transition.

The averaging mechanism can explain the underestimation of the effect of x during the first transition. The models implicit in panels (a) and (b) have subtly different dependent variables: in panel (a) one is modelling the probability that an *individual* passes the transitions, while in panel (b) one models the *average* probability of passing the transitions. The two result in different estimates because the relationship between the unobserved variables and the probabilities is non-linear. This issue is discussed in terms of the sequential logit model by Cameron and Heckman (1998). It also occurs in other non-linear models, and has been discussed by Neuhaus et al. (1991), Allison (1999) and Mood (2010). It is also closely related to the distinction between population average or marginal models on the one hand and mixed effects or subject specific models on the other (Fitzmaurice et al. 2004, chapter 13; Agresti 2002, chapter 12). The averaging of the probabilities can be seen in Table 7.1: for example the probability of passing transition 2 for low x individuals when not controlling for z is (100 $\times$ 0.25 + 200 $\times$ 0.5)/300 = 0.417. The consequence of this is that if we think that equations (7.1) and (7.2) form the true model for the probabilities of passing the transitions, then the true model for the probabilities averaged over z should be represented by equations (7.5) and (7.6), where $\mathrm{E}_z(u)$ is the average of u over z. Instead, the model represented by equations (7.7) and (7.8) are estimated when z is not observed and z is thus left out of the model. The two models are not the equivalent because $\Lambda()$ is a non-linear transformation. Neuhaus and Jewell (1993) give an approximation of how β^*_{11} and β^*_{12} deviate from β_{11} and β_{12}: β^*_{11} and β^*_{12} will be smaller than β_{11} and β_{12}, and the difference between the estimates β^*_{11} and β^*_{12} and the estimates β_{11} and β_{12} will increase when the variances of $\beta_{21}z$ and $\beta_{22}z$ increase and when the probability of passing is closer to 50%.

$$\begin{aligned} \mathrm{E}_z(\Pr[y \in \{B,C\}|x,z]) &= \mathrm{E}_z(\Lambda(\beta_{01} + \beta_{11}x + \beta_{21}z)) \quad (7.5) \\ \mathrm{E}_z(\Pr[y \in \{C\}|x,z,y \in \{B,C\}]) &= \mathrm{E}_z(\Lambda(\beta_{03} + \beta_{12}x + \beta_{22}z)) \quad (7.6) \end{aligned}$$

$$\begin{aligned} \mathrm{E}_z(\Pr[y \in \{B,C\}|x,z]) &= \Lambda(\beta^*_{01} + \beta^*_{11}x) \quad (7.7) \\ \mathrm{E}_z(\Pr[y \in \{C\}|x,z,y \in \{B,C\}]) &= \Lambda(\beta^*_{02} + \beta^*_{12}x) \quad (7.8) \end{aligned}$$

7.3 A sensitivity analysis

The previous section discussed what kind of problems unobserved variables might cause. The difficulty with finding a solution for these problems is that it is obviously challenging to control for something that has not been observed. One possible solution is to perform a sensitivity analysis: specify a number of plausible scenarios concerning the unobserved variables, and estimate the effects within each scenario. The aim of this type of analysis is not to get an empirical estimate of the effect *per se*, but to assess how important assumptions are for the estimated effect and to get a feel for the range of plausible values for the effect. There are many potential problems that could all simultaneously influence the results of an analysis and whose influence could all be investigated using sensitivity analysis. However, to give the analysis focus it is often better to narrow down the scope of the sensitivity analysis by concentrating on a specific subset of potential problems. For example, the aim of the sensitivity analysis proposed in this chapter is to assess the sensitivity to the effect of unobserved heterogeneity through the selection mechanism and averaging mechanism.

A key step in creating such scenarios is to create a set of reasonable scenarios concerning the unobserved variable z. In the example in the previous section, z was assumed to be dichotomous, because that would result in an easy numerical example. When creating the scenarios, it is more useful to think about z as not being a single unobserved variable but as a (weighted) sum of all the unobserved variables. Such a sum of random variables can usually be well approximated by a normal distribution, even if the constituent variables are non-normally distributed. So, it is reasonable to represent the distribution of the composite unobserved variable with a normal distribution. There are two equivalent ways of thinking about the scale of this compound unobserved variable. It is sometimes convenient to think of the resulting variable as being standardized, such that mean is 0 and the standard deviation is 1. This way the 'effect' — call that γ — can be compared with the effects of standardized observed variables to get a feel for the range of reasonable values of this 'effect'. Alternatively, it is possible to think of the composite unobserved variable as just being an unstandardized random variable or error term. In this case, the standard deviation of this random variable is the same as γ. The standardized unobserved variable will be referred to as z, while the unstandardized unobserved variable will be referred to as ε in order to distinguish between the two. The two are related in the following way: $\gamma z = \varepsilon$.

In this chapter I will propose a set of scenarios based on this representation of the unobserved variable. This basic scenario is introduced in equations (7.9) till (7.12). In this example there are two transitions, with the probabilities of passing these transitions influenced by two variables x and z, where z is as defined above. The observed dependent variables are the probabilities of passing the two transitions averaged over z. So by estimating models (7.9) and (7.11), one can recover the true effects of x. To estimate it, all one needs to know is the distribution of γz ($= \varepsilon$) and to integrate over this distribution, as in equations (7.10) and (7.12). The mean of ε will be set at 0 and a standard deviation equal to γ, which is *a priori* fixed in the scenario. Furthermore, it assumes that a person's value on ε will not change over the transitions, implicitly assuming that both the value on z and the effect of z (γ) will not change over the transitions[2].

$$\mathrm{E}_\varepsilon(\Pr[y \in \{B, C\}|x, \varepsilon]) = \mathrm{E}_\varepsilon(\Lambda(\beta_{01} + \beta_{11}x + \underbrace{\gamma z}_{\varepsilon})) \tag{7.9}$$

$$= \int \Lambda(\beta_{01} + \beta_{11}x + \varepsilon) f(\varepsilon) d\varepsilon \tag{7.10}$$

$$\mathrm{E}_\varepsilon(\Pr[y \in \{C\}|x, \varepsilon, y \in \{B, C\}]) = \mathrm{E}_\varepsilon(\Lambda(\beta_{02} + \beta_{12}x + \underbrace{\gamma z}_{\varepsilon})) \tag{7.11}$$

$$= \int \Lambda(\beta_{02} + \beta_{12}x + \varepsilon) f(\varepsilon | y \in \{B, C\}) d\varepsilon \tag{7.12}$$

The effects in each scenario are estimated using maximum likelihood. Referring back to Figure 7.1, the likelihood function for an individual i can be written as equa-

[2]All these assumptions can be relaxed, but relaxing these assumptions will quickly lead to an unmanageable number of scenarios. Moreover, these complications would not contribute to the aim of these scenarios, which assess the sensitivity of estimates to unobserved heterogeneity through the selection mechanism and averaging mechanism.

tion (7.13), that is, the probability of observing someone with value A equals the probability of failing the first transition, the probability of observing someone with value B equals the probability of passing the first transition and failing the second transition, and the probability of observing someone with value C equals the probability of passing both transitions.

$$L_i = \begin{cases} 1 - p_{1i} & \text{if } y_i = A \\ p_{1i} \times (1 - p_{2i}) & \text{if } y_i = B \\ p_{1i} \times p_{2i} & \text{if } y_i = C \end{cases} \tag{7.13}$$

By replacing p_{1i} with equation (7.10) and p_{2i} with equation (7.12), one gets a function that gives the probability of an observation, given the parameters β. This probability can be computed for each observation and the product of these form the probability of observing the data, given a set of parameters. Maximizing this function with respect to the parameters give the maximum likelihood estimates. These estimates include the true effects of the variable of interest x assuming that the model for the unobserved heterogeneity, in particular the standard deviation of ε, is correct.

The difficulty with this likelihood is that there are no closed form solutions for the integrals in equations (7.10) and (7.12). This can be resolved by numerically approximating these integrals using maximum simulated likelihood (Train, 2003). Maximum simulated likelihood uses the fact that the integral is only there to compute a mean probability. This mean can be approximated by drawing at random many values for ε from the distribution of ε, computing the probability of passing a transition assuming that this randomly drawn value is the true value of ε, and then computing the average of these probabilities. This approach can be further refined by realizing that using true random draws is somewhat inefficient as these tend rather to cluster. Increasing the efficiency is important as these integrals need to be computed for each observation, meaning that these simulations need to be repeated for each observation. One can cover the entire distribution with less draws if one can use a more regular sequence of numbers. An example of a more regular sequence of numbers is a Halton (1960) sequence. A Halton sequence will result in a more regular series of quasi-random draws from a uniform distribution. These quasi-random draws can be transformed into quasi-draws from a normal distribution by applying the inverse cumulative normal distribution function. These are then used to compute the average probability of passing the first transition, as is shown in equation (7.14), where m represents the number of draws from the distribution of ε. At subsequent transitions, the distribution of ε is no longer a normal distribution, but conditional on being at risk. The integral over this distribution is computed by drawing ε from a normal distribution as before, but then computing a weighted mean whereby each draw is given a weight equal to the probability of being at risk assuming that that draw was the true ε. In the appendix to this chapter I show that this is a special case of importance sampling (Robert and Casella, 2004, 90–107). This procedure is implemented in the `seqlogit` package (Buis, 2007b) in Stata (StataCorp, 2007), using the facilities for generating Halton sequences discussed by Drukker and Gates (2006). This package is documented in Technical Materials II.

$$\mathrm{E}_\varepsilon(\Pr(y \in \{B, C\}|x, \varepsilon)) \approx \frac{1}{m}\sum_{j=1}^{m} \Lambda(\beta_{01} + \beta_{11}x + \varepsilon_j) \tag{7.14}$$

$$\mathrm{E}_\varepsilon(\Pr(y \in \{C\}|x, \varepsilon, y \in \{B, C\}) \approx \frac{\sum_{j=1}^{m}\left[\Pr(y\in\{B,C\}|x,\varepsilon_j)\Lambda(\beta_{02}+\beta_{12}x+\varepsilon_j)\right]}{\sum_{j=1}^{m}\Pr(y\in\{B,C\}|x,\varepsilon_j)} \tag{7.15}$$

7.4 An example: The effect of family background on educational attainment in the Netherlands

An important application for the sequential logit model is the study of the influence of family background on educational attainment (for recent reviews see: Breen and Jonsson, 2005; Hout and DiPrete, 2006). The potential problems that unobserved variables can cause were recognized from the time that the sequential logit model was introduced in this literature (Mare, 1979, 1980, 1981), but interest in this issue has been revived by the critique from Cameron and Heckman (1998). However, only a limited number of empirical studies have tried to actually account for unobserved heterogeneity (for exceptions see: Mare, 1993; Rijken, 1999; Chevalier and Lanot, 2002; Lauer, 2003; Arends-Kuenning and Duryea, 2006; Colding, 2006; Lucas et al., 2007; Holm and Jæger, 2008). The method proposed in this paper will be illustrated by replicating an analysis that does not control for unobserved heterogeneity by De Graaf and Ganzeboom (1993) and in Chapter 2 of the effect of father's occupational status and education on transition probabilities between educational levels in the Netherlands, and assessing how sensitive the conclusions are to assumptions about unobserved heterogeneity. The original study by De Graaf and Ganzeboom (1993) was part of an influential international comparison of the effect of family background on educational attainment (Shavit and Blossfeld, 1993). It used 10 Dutch surveys that were post-harmonized as part of the International Stratification and Mobility File [ISMF] (Ganzeboom and Treiman, 2009). In Chapter 2 I updated this analysis by using an additional 33 Dutch surveys that have since been added to the ISMF.

Figure 7.2: Simplified model of the Dutch education system

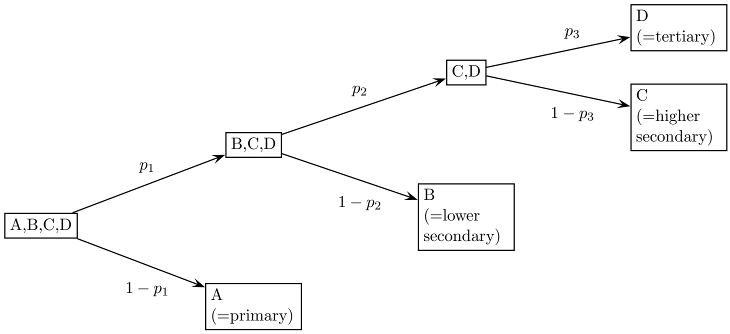

7.4.1 The data

The total of 43 surveys were held between 1958 and 2006. Only male respondents older than 25 are used in the analysis. These surveys contain 35,846 men with valid information on all the variables used in the model. Family background is measured as the father's occupational status and the father's highest achieved level of education. Time was measured by 10-year birth cohorts covering the cohorts that were born between 1891–1980. The main effect of time is added as a set of dummies, while the effects of the family background variables is allowed to change linearly over the cohorts.

The father's occupational status was measured using the International Socio-Economic Index (ISEI) of occupational status (Ganzeboom and Treiman, 2003), which originally ranged between 10 and 90 and was recoded to range between 0 and 8. In concordance with De Graaf and Ganzeboom (1993) and Chapter 2, education of both the father and the respondent were measured in four categories: primary education (LO), lower second secondary education (LBO and MAVO), higher secondary education (HAVO, MBO, and VWO), and tertiary education (HBO and WO). The value of the father's highest achieved level of education was created by giving these educational categories the numerical values 1 till 4. The transitions that were studied by De Graaf and Ganzeboom (1993) and in Chapter 2 are: 1) from primary education or less to a diploma in secondary or tertiary education; 2) from a diploma in lower secondary education to a diploma in higher secondary or tertiary education; 3) from a diploma in higher secondary education to completed tertiary education. These transitions are displayed in Figure 7.2.

7.4.2 The results

The effects of father's occupational status and education are estimated for four scenarios, and the results are represented in the different columns in Table 7.2. The first scenario assumes that the standard deviation of ε is zero, which is a replication of the model used by De Graaf and Ganzeboom (1993) and in Chapter 2. This replication shows three main patterns. First, both father's occupational status and father's education have a positive effect on the probability of passing transitions. Second, this effect decreases over transitions. Third, the effect of father's education decreases over cohorts during all three transitions while the effect of father's occupational status clearly decreases over cohorts for the first transition, but the trend is non-significant negative during the second transition and non-significant positive during the third transition. These patterns are the same as those found by De Graaf and Ganzeboom (1993) and in Chapter 2 with the exception of the significant negative trend in the effect of father's education during the third transition, which was not found to be significant by De Graaf and Ganzeboom (1993).

The remaining three scenarios assume that the standard deviation of ε is .5, 1, and 2. As was discussed before, the standard deviations represent the effects (log odds ratios) if the unobserved variable z is a standardized variable. To put these scenarios into perspective, one can look at the effects of father's occupational status and education when both are standardized in the earliest cohort at the first transition, when the effects are largest. These standardized effects are .823 for father's occupational status and 1.453 for father's education[3]. So, the values .5, 1, and 2 capture a reasonable range of values for the effect of a standardized unobserved variable.

The results from the different scenarios, as presented in the remaining columns of Table 7.2, show that the qualitative conclusions remain unchanged, that is, those effects that were significant remained significant and those that were not significant remained not significant. However, the size of the effects of father's occupational status and education and their trends did change over the scenarios: the effects increased as the amount of unobserved heterogeneity increased, while the trends in the effects over time became more negative, and the decrease in the effects over transitions becomes less pronounced. This is also shown in Figures 7.3 and 7.4. In addition, these figures show that difference between the scenarios decreased over time, indicating that the bias due to unobserved heterogeneity decreased over time. This is particularly strong for the first transition.

In section 7.2 I discussed that unobserved heterogeneity could influence the results through two mechanisms. First, the averaging mechanism is based on the fact

[3]The effects of the unstandardized variables are presented in Table 7.2, and the standard deviation of father's occupational status is 1.55 and the standard deviation of father's education is 1.01.

Table 7.2: Log odds ratios in models for men assuming different degrees of unobserved heterogeneity (the main effects of the cohort dummies and the constant are not displayed)

	sd(ε) = 0	sd(ε) = .5	sd(ε) = 1	sd(ε) = 2
primary v lower secondary				
father's education	1.439	1.496	1.641	2.092
	(11.50)	(11.56)	(11.70)	(12.10)
father's education X cohort	-0.117	-0.124	-0.142	-0.192
	(-4.80)	(-4.96)	(-5.28)	(-5.87)
father's occupation	0.531	0.558	0.628	0.833
	(13.08)	(13.22)	(13.46)	(13.73)
father's occupation X cohort	-0.057	-0.061	-0.070	-0.097
	(-6.34)	(-6.57)	(-7.02)	(-7.60)
lower secondary v higher secondary				
father's education	0.713	0.796	0.995	1.512
	(11.79)	(12.50)	(13.86)	(15.94)
father's education X cohort	-0.026	-0.034	-0.051	-0.092
	(-2.34)	(-2.88)	(-3.88)	(-5.31)
father's occupation	0.294	0.333	0.424	0.655
	(8.10)	(8.73)	(9.96)	(11.90)
father's occupation X cohort	-0.010	-0.014	-0.023	-0.045
	(-1.49)	(-2.01)	(-2.98)	(-4.42)
higher secondary v tertiary				
father's education	0.445	0.539	0.748	1.252
	(7.11)	(8.14)	(10.05)	(12.99)
father's education X cohort	-0.031	-0.039	-0.057	-0.099
	(-2.78)	(-3.34)	(-4.31)	(-5.70)
father's occupation	0.149	0.187	0.275	0.486
	(3.41)	(4.07)	(5.34)	(7.42)
father's occupation X cohort	0.010	0.007	0.001	-0.011
	(1.25)	(0.87)	(0.15)	(-0.97)

(z-values in parentheses)

Figure 7.3: The effect of father's occupational status

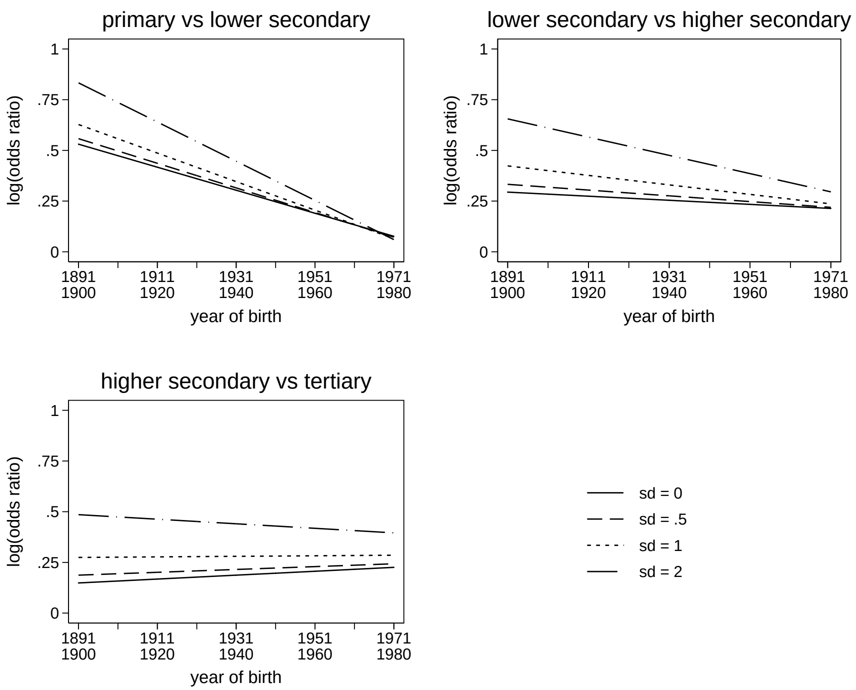

Figure 7.4: The effect of father's education

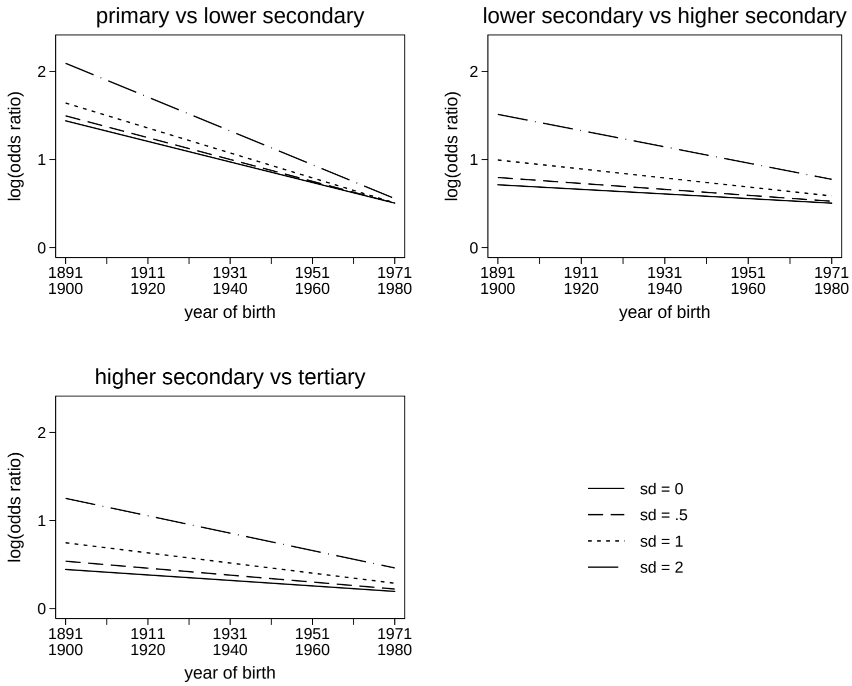

that a model that leaves out the unobserved variables models the average probability of passing the transitions rather than an individual's probability of passing. This will lead to an underestimation of the effect if one leaves the variable out of the model, and this bias will be larger when the variance of the unobserved variable increases and when the probability of passing is closer to 50% (Neuhaus and Jewell, 1993). Second, the selection mechanism is based on the fact that after the first transition the unobserved variable becomes correlated with the observed variables. This means that at later transitions, leaving the unobserved variable out of the model will result in omitted variable bias, even if the unobserved variable was not a confounding variable at the first transition. A key element in both mechanisms is the distribution of the unobserved variable. Table 7.3 shows how the distribution of the unobserved variable changes over the transitions for the different scenarios for men born between 1931 and 1940 (the largest cohort in the data). The first row shows the proportion of respondents at risk of passing this transition, which indicates how selective a transition is. The second and third set of rows shows for each scenario and transition the correlation between the unobserved variable and father's occupational status, and between the unobserved variable and father's education, respectively. This correlation captures the selection mechanism. At the first transition this correlation is by definition 0, but at later transitions it becomes negative, leading to an underestimation of the effect of father's occupational status and education at the later transitions. The correlation becomes larger at later transitions and when the variance of the unobserved variable increases. The correlation between ε and father's education is stronger than the correlation between ε and the father's occupational status. The reason for this is that the correlation is the result of the selection on all the variables at the earlier transitions, and the selection on father's education is stronger than the selection on father's occupational status (the standardized coefficients for the main effects are, as was mentioned before, 1.453 and .823 respectively). The fourth set of rows shows that the variance of the unobserved variable, which plays a key role in the averaging mechanism, and which decreases somewhat over the transitions, but not much. The fifth set of rows shows that the respondents score higher than average on the unobserved variable at the higher transition.

Table 7.3: Changes in the distribution of the unobserved variable over the transitions for men born between 1931 and 1940

		primary v lower secondary	lower secondary v higher secondary	higher secondary v tertiary
Pr(at risk)		1	.837	.487
corr(ε,	sd(ε) = 0	0	0	0
father's occupation)	sd(ε) = 0.5	0	-0.028	-0.070
	sd(ε) = 1	0	-0.051	-0.124
	sd(ε) = 2	0	-0.081	-0.187
corr(ε,	sd(ε) = 0	0	0	0
father's education)	sd(ε) = 0.5	0	-0.048	-0.111
	sd(ε) = 1	0	-0.087	-0.193
	sd(ε) = 2	0	-0.134	-0.282
sd(ε)	sd(ε) = 0	0	0	0
	sd(ε) = 0.5	0.5	0.492	0.480
	sd(ε) = 1	1	0.950	0.883
	sd(ε) = 2	2	1.764	1.531
mean(ε)	sd(ε) = 0	0	0	0
	sd(ε) = 0.5	0	0.038	0.132
	sd(ε) = 1	0	0.143	0.460
	sd(ε) = 2	0	0.460	1.313

Table 7.3 gives an idea of the distribution of the unobserved variable at one point in time, but it cannot explain why this bias changed over time, as was shown in Figures 7.3 and 7.4. The way unobserved heterogeneity influences the results is a function of the proportion of respondents that are at risk at each transition and these have changed considerably over time as is shown in Figure 7.5. As in most other countries, younger cohorts will on average receive more education than the older cohorts, so the proportion of respondents at risk increases over time. Figure 7.5 also explains why the bias in the first transition decreases. The bias in the first transition is due to the averaging mechanism, and the bias due to the averaging mechanism will decrease when the probability of passing approaches 1 (or 0) (Neuhaus and Jewell, 1993). The proportion of respondents that passed the first transition is the proportion at risk of passing the second transition. Figure 7.5 shows that this proportion increased dramatically and is now virtually 1, thus leading to a reduction in the size of the bias. Figure 7.6 showed how the correlation between the father's occupational status and education and the unobserved variable changed over time. It shows that this correlation strongly decreased over time as the higher transitions became less selective, and thus that the bias due to the selection mechanism decreased over time. Figure 7.7 shows that the standard deviation of the unobserved variable hardly changes over time. Figure 7.8 shows how the mean of the unobserved variable decreases at each subsequent transition and how these transitions have become less selective over time.

In summary, this replication showed that the qualitative conclusions from De Graaf and Ganzeboom (1993) and Chapter 2 are largely robust against assumptions on unobserved heterogeneity. However, the scenarios also showed that the size of the effects and the trends are likely to have been underestimated because the original sequential logit models estimated the effect on the average probability of passing rather than on an individual's probability of passing, and because the unobserved variable and the observed variables became negatively correlated at the higher transitions.

Figure 7.5: The proportion of respondents at risk of passing each transition

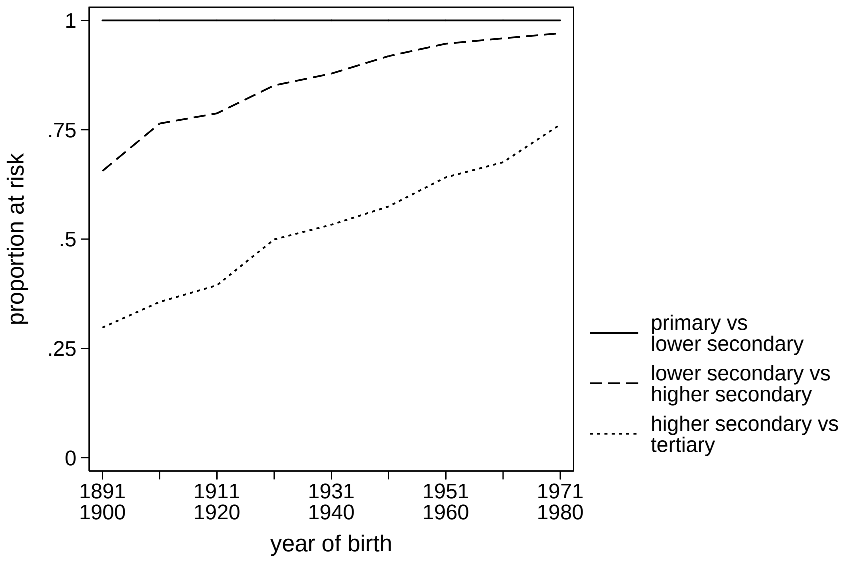

Figure 7.6: The correlation between the unobserved variable and father's occupational status and father's education

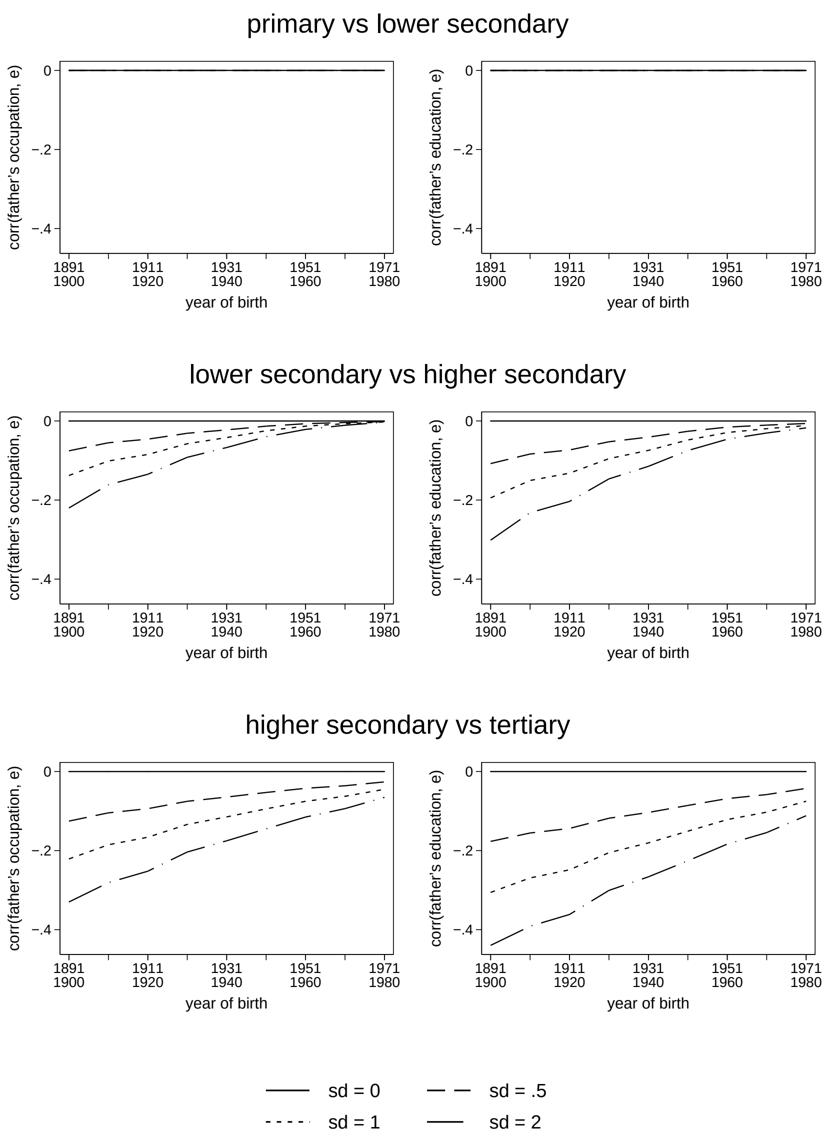

Figure 7.7: The standard deviation of the unobserved variable

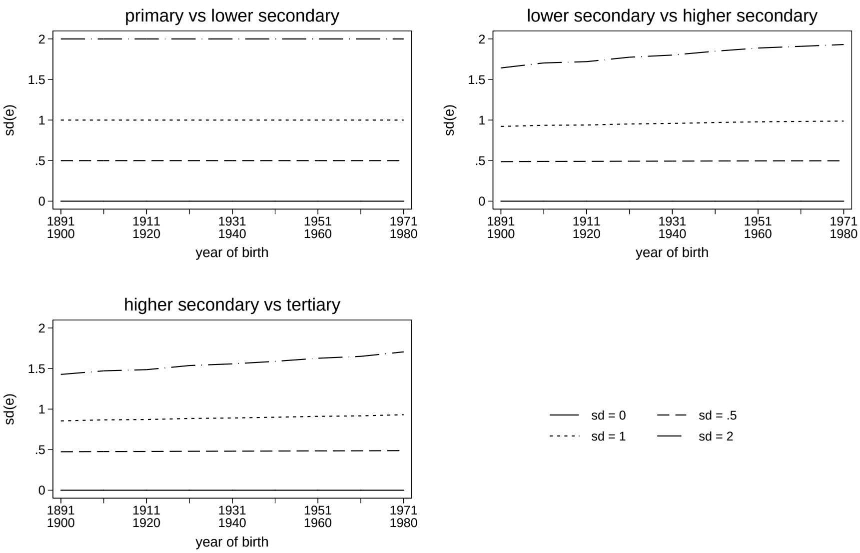

Figure 7.8: The mean of the unobserved variable

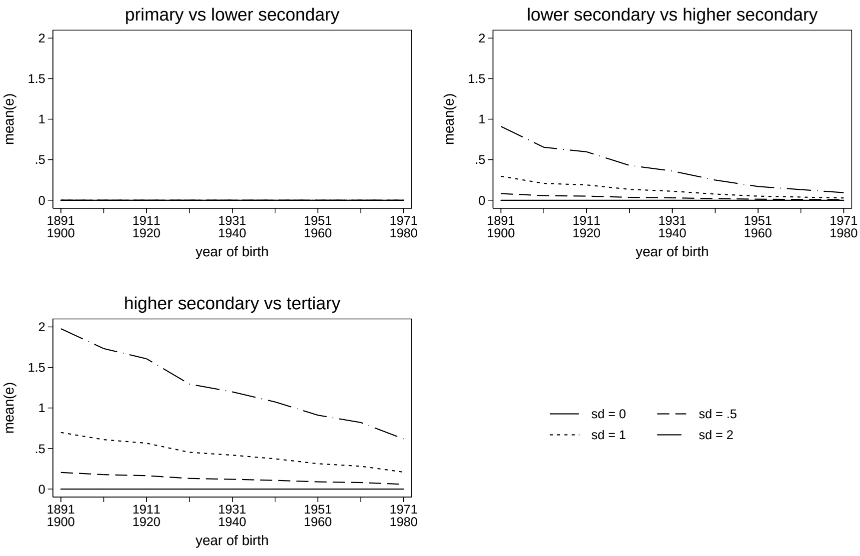

7.5 Conclusion and discussion

The aim of this chapter is to present a sensitivity analysis that can be used to investigate the consequences of unobserved variables in a sequential logit model, and in particular the consequences of leaving a non-confounding variable out of a sequential logit model as discussed by Cameron and Heckman (1998). The bias that these unobserved variables cause are shown to be the result of two mechanisms: first, the averaging mechanism is based on the fact that when a variable is left out of the model, one models the probability of passing the transitions averaged over the variable that is left out. As a consequence, just leaving the unobserved variable out of the model will lead to estimates of effects of the observed variables on the probability of passing the transitions averaged over the unobserved variables rather than the effects on the individual's probability of passing. These two are different because the unobserved variable is related to the probabilities through a non-linear function. Second, the selection mechanism is based on the fact that a variable that is not a confounding variable at the first transition is likely to become a confounding variable at the later transitions. The reason for this is that the process of selection at the earlier transitions will introduce correlation between the observed and unobserved variables.

The method proposed in this chapter to investigate the consequences of unobserved heterogeneity is to perform a sensitivity analysis by specifying a set of scenarios regarding the extent of unobserved heterogeneity, and estimating the effects of the observed variables given those scenarios. This will not give an empirical estimate of the effects of interest, but does give an idea about the sensitivity of the estimates to assumptions about unobserved heterogeneity, and direction of the bias, the size of the bias, and the range of likely values of the effect. The scenarios that have been proposed in this chapter are constructed in the following way: the unobserved variable is normally distributed, for each individual the value of this unobserved variable is assumed to remain constant over the transitions, and the effect of the unobserved variable also remains constant over the transitions. The scenarios differ from one another with respect to the variance of the unobserved variable. This way one can compare what happens to the effects of the observed variables when there is a small, medium, and large amount of unobserved heterogeneity. Moreover, it is possible to recover the distribution of the unobserved variable at the later transitions. This makes it possible to see how, in each scenario, the correlation between the observed and unobserved variables change over the transitions, and/or over a third variable, for example time. The effects of the observed variables within each scenario are estimated by maximum likelihood. The likelihood is defined by integrating over the unobserved variable, which is done using Maximum Simulated Likelihood (Train, 2003).

This method was illustrated by replicating a study by De Graaf and Ganzeboom

(1993) and in Chapter 2 on the effect of the father's occupational status and education on the offspring's educational attainment. The proposed analysis showed that the results of statistical tests were rather robust to changes in the assumptions about unobserved heterogeneity, but that the effects of both the father's occupational status and the father's education were likely to be underestimated, as these effects were stronger in scenarios with more unobserved heterogeneity. Scenarios with more unobserved heterogeneity also resulted in a stronger downward trend over time in the effect of father's occupational status and education. The decrease in the effect of father's occupational status and education over transitions became less in scenarios with more unobserved heterogeneity. This indicates that the commonly found pattern of decreasing effects of family background variables over transitions is at least in part due to unobserved heterogeneity.

This chapter can be seen as part of a larger effort aimed at obtaining an empirical estimate of the causal effect of family background while controlling for unobserved variation between individuals (Mare, 1993, 1994; Cameron and Heckman, 1998; Lucas et al., 2007; Holm and Jæger, 2008). The challenge of this literature is that it tries to solve an unsolvable problem, since obtaining an empirical estimate is by definition incompatible with controlling for unobserved variation. On the one hand this means that it is very unlikely that a single study can build a completely convincing empirical argument for such an effect. On the other hand, that does not mean that estimates obtained in these studies contain no information whatsoever. The key is that each of these methods exploits different parts of the data to get an approximation of the effect. For example, Mare (1993, 1994) uses the nesting of individuals within families, Lucas et al. (2007) and Holm and Jæger (2008) use the presence of instrumental variables, and Mare (2006) uses the strong assumption that all changes in the effect of the explanatory variables over transitions is due to unobserved heterogeneity. In the long run, these differences in strategy can be used to get a plausible range for the causal effect of family background by collecting a sufficient body of evidence using these different methods, followed by an analysis of how the differences in strategy has led to the differences and similarities in the conclusions of these studies.

Appendix: Sampling from the distribution of ε conditional on having passed the previous transitions

One method of sampling from a distribution is importance sampling (Robert and Casella, 2004, 90–107). This appendix will show that the method used in this chapter is a special case of importance sampling. The idea behind importance sampling is that instead of sampling from the distribution of interest $f(\varepsilon)$ one draws samples from another distribution $g(\varepsilon)$, and compute the mean by weighting each draw by $\frac{f(\varepsilon_j)}{g(\varepsilon_j)}$, so one could approximate $\mathrm{E}_\varepsilon[\Lambda(\beta_{02} + \beta_{12}x + \varepsilon)]$ with equation (7.16).

$$\mathrm{E}_\varepsilon[\Lambda(\beta_{02} + \beta_{12}x + \varepsilon)] \approx \frac{1}{m}\sum_{j=1}^{m}\frac{f(\varepsilon_j)}{g(\varepsilon_j)}\Lambda(\beta_{02} + \beta_{12}x + \varepsilon) \tag{7.16}$$

In this chapter the distribution of interest is the distribution conditional on being at risk, while the other distribution is the distribution not conditional on being at risk. These distributions are independent of x, so the conditioning on x in equation (7.17) is superfluous, but this will prove useful later on.

$$\mathrm{E}_\varepsilon[\Lambda(\beta_{02} + \beta_{12}x + \varepsilon)] \approx \frac{1}{m}\sum_{j=1}^{m}\frac{f(\varepsilon_j|x, y \in \{B, C\})}{f(\varepsilon_j|x)}\Lambda(\beta_{02} + \beta_{12}x + \varepsilon) \tag{7.17}$$

Instead of using equation (7.17) directly, the integral is computed using equation (7.18). The aim of this appendix is to show that these two are equivalent.

$$\mathrm{E}_\varepsilon[\Lambda(\beta_{02} + \beta_{12}x + \varepsilon)] \approx \frac{\sum_{j=1}^{m}\left[\Pr(y \in \{B, C\}|x, \varepsilon_j)\Lambda(\beta_{02} + \beta_{12}x + \varepsilon)\right]}{\sum_{j=1}^{m}\Pr(y \in \{B, C\}|x, \varepsilon_j)} \tag{7.18}$$

The denominator of equation (7.18) can be rewritten as in equation (7.19), which leads to equation (7.20)

$$\begin{aligned}\sum_{j=1}^{m}\Pr(y \in \{B, C\}|x, \varepsilon_j) &= m\frac{\sum_{j=1}^{m}\Pr(y \in \{B, C\}|x, \varepsilon_j)}{m} \\ &\approx m\Pr(y \in \{B, C\}|x)\end{aligned} \tag{7.19}$$

$$\mathrm{E}_{\varepsilon}[\Lambda(\beta_{02} + \beta_{12}x + \varepsilon)] \approx \frac{1}{m}\sum_{j=1}^{m}\frac{\Pr(y \in \{B, C\}|x, \varepsilon_j)}{\Pr(y \in \{B, C\}|x)}\Lambda(\beta_{02} + \beta_{12}x + \varepsilon) \quad (7.20)$$

Comparing equations (7.17) and (7.20) indicates that the problem can be simplified to showing that equation (7.21) is true.

$$\frac{f(\varepsilon_j|x, y \in \{B, C\})}{f(\varepsilon_j|x)} = \frac{\Pr(y \in \{B, C\}|x, \varepsilon_j)}{\Pr(y \in \{B, C\}|x)} \quad (7.21)$$

Equation (7.21) can be rewritten as equation (7.22). Using Bayes' theorem, equation (7.22) can be rewritten as equation (7.23). Equation (7.23) is true, thus showing that equations (7.17) and (7.18) are equivalent. Notice, however, that this is based on the approximation in equation (7.19), which will get better as the number of samples m increases.

$$f(\varepsilon_j|x, y \in \{B, C\})\Pr(y \in \{B, C\}|x) = \Pr(y \in \{B, C\}|x, \varepsilon_j)f(\varepsilon_j|x) \quad (7.22)$$

$$f(\varepsilon_j \cap y \in \{B, C\}|x) = f(y \in \{B, C\} \cap \varepsilon_j|x) \quad (7.23)$$

Chapter 8

Conclusions and discussion

In this dissertation I have investigated the changing association between family background and educational attainment in the Netherlands during the 20th century. This association is a measure of the inequality in access to education, as it indicates the extent to which persons with a more privileged background are more likely to attain a higher level of education than persons with a less privileged background. This inequality in access to education is not only important to investigate because education is a valuable and scarce resource in its own right, but also because it influences future success in other domains of life, like work, family formation, and health. The research literature on the inequality of access to education has a long history (Hout and DiPrete, 2006; Breen and Jonsson, 2007). This dissertation contributed to this literature by studying the following aspects of inequality in access to education: 1) the inequality in the outcome of the process of attaining education, 2) the inequality during the process of attaining education, as well as the relationship between these two types of inequalities. I have labelled these two types of inequality Inequality of Educational Outcome (IEOut) and Inequality of Educational Opportunity (IEOpp) respectively. The overarching research question that guided the individual studies that make up this dissertation has been: “To what extent, how, and when has a trend toward less inequality in educational opportunities and in educational outcomes of persons from different family backgrounds occurred in the Netherlands?”

As a point of departure I replicated in Chapter 2 a study by De Graaf and Ganzeboom (1993) using more, and more recent data. This replication served as a benchmark, as it represents what can be learned from the most recent data using ‘default’ methods. The remaining chapters consisted of applying new methods that improved on these ‘default’ methods. Chapters 3, 4, and 5, showed three ways of improving the estimates of IEOut: In Chapter 3 a scale of education was empirically estimated to replace the *a priori* scale that has been used in the ‘default’ method. In Chapter 4 the trend in IEOut was estimated using a local polynomial curve which is more flexible than the quadratic curve and more powerful than the discrete curve that have been used in the ‘default’ approach. In Chapter 5 a new method was introduced for testing whether the relative differences in effects of occupational status and education of the father and the mother on the offspring’s educational attainment have changed over time. Chapter 6 showed a new way of relating IEOpps to IEOut, which also turned out to provide a meaningful way of analyzing the effect of educational expansion on IEOut. Chapter 7 showed a way of improving the estimates of the IEOpps, by proposing a sensitivity analysis to assess the potential impact of unobserved variables on the results.

The conclusions from all these chapters will first be discussed in detail, and are then summarized by answering the overarching research question. Finally, some shortcomings of these studies are discussed together with some recommendations for future research.

8.1 Conclusions

8.1.1 A replication

The dissertation started with a replication of the study by De Graaf and Ganzeboom (1993), which was the Dutch contribution to an influential international comparative project by Shavit and Blossfeld (1993). The role of the replication in this dissertation is to create a point of reference in terms of the estimated trend in inequality of access to education using 'default' methods. De Graaf and Ganzeboom (1993) studied IEOpp and IEOut, which both play a prominent role in this dissertation. Moreover, the data used in this dissertation is an extension of the data used by De Graaf and Ganzeboom (1993). They used data from ten cross-sectional surveys that were post-harmonized and then stacked to form a single dataset. Ganzeboom and Treiman (2009) have since extended this data as part of the International Stratification and Mobility File (ISMF) such that the Dutch part of this file now contains information from 54 surveys. It is this data that has been used throughout this dissertation.

The main finding of this replication is that despite the fact that this replication used more than five times as many respondents (69,868 versus 11,244 respondents) and covered 20 additional years (1891–1980 versus 1891–1960), the results remained largely unchanged. Using default methods on the extended dataset the following trends in IEOpp and IEOut were found for the Netherlands: a significant negative trend in IEOut and a significant negative trend in IEOpp for the transition whether or not to continue after primary education; mixed evidence for the trend in IEOpp for the choice of track during secondary education; and no trend and in some cases a positive trend for the transition whether or not to finish tertiary education. Moreover, these trends are mostly found to be linear.

8.1.2 IEOut: operationalizing education, the trend, and family background

Chapters 3, 4, and 5 of this dissertation proposed three ways of improving on the 'default method' of estimating the trend in IEOut.

Chapter 3 focussed on the values assigned to each of the educational categories. These values are necessary in order to estimate IEOut. Following De Graaf and Ganzeboom (1993), the replication in Chapter 2 assigned values to educational categories by distinguishing between four educational categories (primary, lower secondary, higher secondary, and tertiary education) and assigning them the values 1 to 4. A major advantage of this method is that it is easy to apply, all that is necessary is a rank order of the educational categories. A disadvantage is that this method implicitly assumes that the distances between successive educational categories are all equal. An often-used alternative approach is to assign each category a value equal to the number of years it would take a 'standard' student to complete that category. An advantage of this method is that these values can easily be derived for most educational systems from (semi-)official documents. However, it conflates two distinct concepts: the duration and the value. As a result, the rank ordering of educational categories based on these standard durations sometimes does not correspond to the rank ordering based on a priori knowledge about the values. This is the case in the Netherlands for higher secondary vocational education [MBO], which has led Ganzeboom and Treiman (2009) to apply an *ad hoc* adjustment to their scale values when creating their scale for the ISMF. Another potential problem with these *a priori* scales of education is that they assume that the values of the educational categories have remained constant over time, while there are two plausible mechanisms through which the value of an educational category could change over time: educational reform, which can mean that an educational category before and after a reform should be treated as two different categories, and changes in the supply of highly schooled labor relative to the demand for highly schooled labor, so called 'diploma inflation'.

Chapter 3 improved these standard ways of assigning values to the educational categories by empirically estimating a scale of education. This scale of education is estimated such that it is optimal for predicting occupational status, using a model with parametricaly weighted covariates proposed by Yamaguchi (2002). The resulting scale largely corresponds with the *a priori* scale by Ganzeboom and Treiman (2009). The major deviation from the *a priori* scale is that the *a priori* scale overrates the value of lower secondary vocational education [LBO], which means that respondents with LBO had, on average, lower status occupations than was predicted using the *a priori* scale. The resulting scale also showed that there is little evidence that the values changed over time, as measured by the year in which the survey was held. The time at which the survey was held was used as a proxy for the time when the respondents held their job. As a consequence, the lack of change over time is an indication that changes in the labor market during the period that was studied (1958 to 2006) had little effect on the relative distances between educational categories. However, the values of two educational categories did differ when comparing cohorts that were in education before and after a major educational reform in the Netherlands, the “Mammoet Wet” or “Mammoth Law”, which was implemented in 1968. The categories that were influenced by it were lower general secondary education [MAVO] and higher professional education [HBO]. The change in the value of MAVO was to be expected, as this diploma changed from a level that prepared for the labor market to a level that prepared for a subsequent level of education (MBO). A possible reason for the change in the value in HBO could be that it became accessible from higher general secondary education [HAVO].

Showing the consequences of using this new scale rather than the *a priori* scale for the estimates of IEOut was one of the subsidiary aims of Chapter 4. The main aim of this chapter was to assess whether or not the trend in IEOut has changed over time. Past research had found a steady negative trend in IEOut, and found no evidence for any non-linearity in this trend. However, it is implausible that this linear trend will continue as this would eventually result in a negative association between family background and educational attainment. So, at one point in time the negative trend in IEOut will have to slow down, and the aim of Chapter 4 was to try to detect this declaration of the trend. This chapter hypothesized that the lack of evidence for a non-linear trend in the default approach was due to the methods used in testing for non-linearities: these methods either estimated a non-linear trend using a quadratic trend, which could be not flexible enough to adequately detect any non-linearities in the trend, or as a discrete trend model, which could be too flexible and thus not powerful enough. The alternative proposed in this chapter was to represent the trend as a local polynomial curve, which is more flexible than a quadratic curve but more powerful than a discrete curve.

This chapter did find evidence that the trend has been non-linear, but did not find the expected deceleration in the decreasing trend in IEOut. A period of negative trend was found for both men (1941–1960) and women (1952–1977), which was preceded by a period of significantly accelerating trend (1935–1944 for men and 1949–1952 for women). There is some evidence — only for men — that the negative trend decelerated prior to becoming insignificant, but this deceleration is not (yet) significant. There is no indication that the negative trend for women decelerated prior to becoming insignificant, indicating that the lack of significance of the negative trend in the youngest cohorts has more to do with lack of power than with a lack of negative trend. Surprisingly, the trend did not show any effect of a major educational reform, the

'Mammoet Wet' or 'Mammoth Law', which was aimed at reducing IEOut and was implemented in 1968.

A subsidiary aim of this chapter was to assess the robustness of these conclusions to three potential sources of error: different scales of education, differences in quality of the data across surveys, and missing data. Using the scale of education estimated in chapter 3 rather than the *a priori* scale by Ganzeboom and Treiman (2009) led to a slightly more stable trend, but did not qualitatively change the conclusions. Controlling for differences between surveys led to a decrease in trend for the earliest cohorts, while using multiple imputation to control for missing values did not influence the results.

In Chapter 5 I assessed which resource — occupational status or education — and which parent — the father or the mother, the highest educated/status parent or the lowest educated/status parent, or the parent with the same sex as the respondent or the parent with a different sex to the respondent — contributed most to the offspring's educational attainment. The results indicate that the distinction between highest and lowest status parent is the main distinction between the parents, rather than the distinction between fathers and mothers or the distinction between the parent with the same sex as the respondent or a different sex to the respondent. There is also moderate evidence that occupational status is more important than parental education. I also found that the mother being a homemaker had a negative effect on the educational attainment of the offspring if the mother has little education and the father has a low status job, but that this effect becomes positive when the mother is well-educated or when the father has a high status job.

In this chapter I also investigated whether the relative contributions of each of these resources changed over time. I expected the value of the contributions of the mother's resources to have increased over time relative to the values of the resources contributed by the father due to changes in the roles of men and women in society during the period studied (1939 till 1991). I also expected the value of occupational status to decline as it is more closely related to the economic resources available in the family, and economic constraints have become less likely to limit the educational choices as almost everybody has become wealthier and education has become more heavily subsidised during the period studied. In order to test these hypotheses, I used a model with parametricaly weighted covariates proposed by Yamaguchi (2002), which estimates the model under the null hypothesis that the relative contributions of these resources have remained unchanged over time. Contrary to my expectations, this hypothesis could not be rejected.

8.1.3 Combining IEOpp and IEOut

When investigating inequality in access to education, it is useful to distinguish between inequality during the process of attaining education (the IEOpp) and the inequality in the final outcome of that process (the IEOut). It is also useful to recognize that IEOpp and IEOut provide complementary information; a discussion of the process of attaining education can be meaningfully supplemented by a discussion of the outcome of that process and *vice versa*. In order to make the best use of this complementarity between IEOpp and IEOut, one needs to move beyond separate discussions of IEOpp and IEOut and towards an integrated discussion of the two. Chapter 6 proposed a new method that makes such an integrated discussion possible. This method starts with the standard model for estimating IEOpps, the sequential logit model as proposed by Mare (1981), which estimates the effect of family background on the probabilities of passing from one level of education to the next. It then shows that this model implies a decomposition of IEOut as a weighted sum of the IEOpps, such that the weights assigned to each transition between levels of education are the product of three elements: 1) the proportion of respondents at risk of passing that transition, which means that a transition receives more weight when more people are affected by it; 2) the variance of the indicator variable indicating whether or not respondents passed that transition, which means that less weight is given to transition where virtually everybody fails or virtually everybody passes; and 3) the expected increase in highest achieved level of education due to passing that transition, which means that a transition receives more weight when passing it is more profitable. This makes it possible to supplement the IEOpps with estimates of how relevant these IEOpps are for IEOut. Moreover, it provides a substantively interpretable mechanism through which educational expansion can influence educational inequality, as educational expansion influences the probabilities of passing the educational transitions, which influence the weights, which in turn lead to changes in IEOut.

When applying this methodology to the Netherlands, I distinguished four transitions: continue or not after primary education, taking the vocational track versus the academic track, continue to higher secondary vocational education given that a respondent entered the vocational track, continue to university given that a respondent entered the academic track. I found that the latter two transitions not only have low IEOpps, which was already known, but they also have low weights. These low weights were primarily due to the relatively low proportion of respondents that are at risk of passing these higher transitions compared to the lower transitions. By contrast, educational expansion had a big influence on the first two transitions. The first transition, whether or not to continue after primary education, started out as the main source of IEOut, but declined quickly as passing this transition became almost universal. The second transition, whether to enter the vocational track [LBO and MAVO] or academic track [HAVO and VWO], strongly increased in importance as the percentage of people passing that transition increased to about 50%, which resulted in an increase in the variance of the dependent variable, and as more and more people became at risk of passing this transition.

8.1.4 IEOpp: the influence of unobserved variables

The standard model for estimating IEOpps, the sequential logit model, has been subject to an influential critique by Cameron and Heckman (1998). They argue that, like any other model, a sequential logit model cannot include all variables that influence the dependent variable. However, leaving these variables out will influence the results, even if these variables are not confounding variables. These so-called unobserved variables influence the results through two mechanisms. First, the 'averaging mechanism' is based on the fact that when a variable is left out of the model, one models the probability of passing the transitions averaged over the variables that are left out. As a consequence, leaving the unobserved variable out of the model will lead to estimates of effects of the observed variables on the *average* probability of passing within groups defined by the observed variables rather than the effects on the *individual's* probability of passing. These two are different in non-linear models like logistic regression because the unobserved variables are related to the probabilities through a non-linear function. Second, the 'selection mechanism' is based on the fact that a variable that is not a confounding variable at the first transition is likely to become a confounding variable at later transitions. The reason for this is that the process of selection at the earlier transitions will introduce correlation between the observed and unobserved variables at the later transitions.

This suggests that one needs to control for these unobserved variables, but it is by definition impossible to get an empirical estimate that is controlled for variables that have not been observed. However, it is possible to create a scenario, by specifying assumptions about the unobserved variables, and estimating the effects within that scenario. There are roughly two ways in which these scenarios can be used. First, one can try to put as much empirical information as possible into these scenarios. For example Mare (1993, 1994) uses the similarity between siblings to capture the unobserved variables on the family level. Alternatively, one can use a set of scenarios to assess the sensitivity of the estimates to the assumptions. Chapter 7 is an example of this latter approach as it proposed a set of scenarios that is useful for such a sensitivity analysis and a method for estimating the effects within these scenarios. This method was illustrated by replicating the analysis in Chapter 2, showing that the results of statistical tests were robust to changes in the assumptions about unobserved heterogeneity, but that the effects of both father's occupational status and father's education were likely to be underestimated, as these effects were stronger in scenarios with more unobserved heterogeneity. Scenarios with more unobserved heterogeneity also resulted in a stronger downward trend over time in the effect of father's occupational status and education, indicating that the trend in the effects of parental background variables across cohorts is also likely to be underestimated. However, the effect of father's occupational status and education decrease less over transitions in scenarios with more unobserved heterogeneity. This indicates that the commonly found pattern of decreasing effects of family background variables over transitions is at least in part due to unobserved heterogeneity.

8.1.5 Summary

These conclusions can be summarized by explicitly answering the overarching research question: "To what extent, how, and when has a trend toward less inequality in educational opportunities and in educational outcomes between persons from different family backgrounds occurred in the Netherlands?" The answers to this question can be broken up into the following elements:

IEOut

- The trend in IEOut was shown to have decreased during the third quarter of the 20$^{\text{th}}$ century, during which time it approximately halved. This negative trend was preceded by an acceleration of the trend, and there is some indication that IEOut was initially increasing. The non-linearity of this trend is a new finding, as previous studies failed to reject the hypothesis of a linear declining trend.
- The 'Mammoet Wet', a major educational reform in the Netherlands implemented in 1968, did not have a noticeable influence on IEOut.
- An improved scale for the educational categories was created in this dissertation, but this was found to have little effect on the estimated trend in IEOut.
- The relative contributions of the education and occupational status of the father and the mother to the respondent's educational attainment were found to have remained constant over cohorts.

IEOut and IEOpp

- The main driving force behind the trend in IEOut turned out to be the major shift in which transition between educational levels contributed most to IEOut. Initially, the transition between whether or not to continue in education after finishing primary education was the main contributor to IEOut. However, the contribution of this transition quickly declined as passing this transitions became almost universal. At the same time the contribution of the second transition between entering the academic or vocational track increased in importance as more people became at risk of passing that transition and as the number of people entering the academic track and vocational track became more evenly balanced. This shift between the transitions resulted in both the initial increase in IEOut, as the decline of the contribution from the first transition was more than compensated by the increase of the contribution from the second transition, and the subsequent decline in IEOut, as the less unequal second transition replaced the more unequal first transition as the dominant source of inequality.

IEOpp

- At the lowest transitions a declining linear trend in the IEOpps over time was found, while at the higher transitions the evidence became mixed with negative, insignificant, and even positive trends.
- The IEOpps at the lower transitions were higher than the IEOpps at the higher transitions.
- A sensitivity analysis showed that qualitative conclusions are robust, but that both the size of the IEOpps and the size of the trend are likely to be underestimated when the unobserved variables are not accounted for.

8.2 Discussion

What all chapters in this dissertation have in common is that they used data from the International Stratification and Mobility File [ISMF] (Ganzeboom and Treiman, 2009). As a consequence, all these chapters share the strengths and weaknesses associated with this source of data. One of these weaknesses is that the ISMF contains data from surveys of differing quality. Chapter 4 found that controlling for differences between surveys did have a moderate effect on the estimated trend in IEOut. Future research could extend on this finding by also modelling the effects of survey characteristics, thus gaining more insight into the way survey quality influences the substantive conclusions that can be drawn from it. This would turn the variation between the surveys present in the ISMF from a potential weakness into a strength, as this variation can then be used to control for characteristics of the survey in ways that are impossible when analyzing surveys separately or only analyzing surveys with certain (high quality) characteristics.

Another potential weakness is the way time is measured using so-called synthetic cohorts, that is, cohorts that are observed in a cross-sectional survey. These synthetic cohorts are used to estimate the trend in IEOpp and IEOut, and thus play a key role in this dissertation. The key advantage of using synthetic cohorts is that it makes it possible to study a long period of time using a large amount of data. However, there are also problems associated with the use of synthetic cohorts. The first problem is that a synthetic cohort is not a proper sample from the population of people born in a certain year, but a sample from the population of people born in a certain year *and* who are still alive and living in the Netherlands at the time the survey was held. This can be a problem for cohorts that are very old when the survey was held because in these cohorts higher-educated respondents are likely to be over-represented, as higher-educated persons are more likely to live longer. Such a selection on the dependent variable can bias the results (Breen, 1996). This was partly solved in most chapters by only using respondents younger than 65 years old[1]. This way, not enough people will have died for this to have become a problem. The second problem with synthetic cohorts is that education happens over a period of time, so it is not exactly clear which historical period is represented by a cohort. A reasonable choice is to look at the time when the respondent was 12, as in the Netherlands that is the age at which people make the most important decision in their educational career, but any such choice will necessarily be an approximation. This is particularly relevant when studying the consequences of a policy change, as synthetic cohorts will only approximately classify the respondents as being affected or not affected by the policy change.

Another difficulty with the use of cross-sectional surveys like the ISMF is that they do not directly measure which transitions a respondent passed. The transitions a respondent has made are reconstructed based on the respondent's highest achieved level of education and a simplified model of the educational system. In particular, in order to be able to reconstruct a respondent's educational career, such a model must impose that a respondent can only reach a certain level of education through one route. This is a limitation, especially within educational systems consisting of multiple tracks, as it precludes the study of indirect paths through the educational system. This can be of substantive interest as these indirect paths represent 'second chances' open to respondents after they have chosen/been placed in a certain track. As a consequence, the 'synthetic educational careers' in the ISMF preclude the study of the question concerning who benefits most from these second chances: the people

[1] Exceptions are chapters 2 and 7, which replicate the study by De Graaf and Ganzeboom (1993).

with fewer family resources, who were initially disproportionately placed in the lower tracks, or the people with more family resources, who are better capable of making the best use of any loophole in the system.

A way to avoid problems with synthetic cohorts and synthetic educational careers is to use data for which the time at which events took place and the educational career are directly measured. This type of information is available in panel data, where students are followed during their educational career, or in cross-sectional surveys where respondents are asked to retrospectively reconstruct their educational career. However, this does not mean that these sources of data are uniformly better than cross-sectional surveys that only asked for the highest achieved level of education. It is actually striking how much the strengths and weaknesses of these different types of data complement one another. An analysis of panel data and retrospective career data can add to an analysis of highest achieved level data as the panel data and retrospective career data have directly observed time and educational careers. An analysis of the final stages of the educational process and the outcome of the educational process is difficult to make in the panel studies due to attrition, but neither the retrospective career data nor the highest achieved level data suffer from this problem. The available panel studies contain data on only a few cohorts, making it difficult to get a detailed description of changes over time, while both the retrospective career data and the highest achieved level data contain information on many cohorts. However, the retrospective career data contain data on relatively few respondents, meaning that each cohort consists of a small number of respondents. The panel data contain few cohorts, but each cohort contains many respondents. The highest achieved level data tend to contain many cohorts, and each cohort consists of many respondents. The retrospective career data can suffer from the fact that its information is based on what a respondent can remember of events that, for some cohorts, occurred many years previously. The panel data do not suffer from this as the data is collected shortly after the events occur, while the highest achieved level data collects information on the highest achieved level of education, which is much more salient and easier to remember than the entire educational career. Future research could make real progress if it were to exploit these complementarities between the data sources rather than continuing to use them separately.

On a more general level, a discussion of this dissertation needs to confront its rather specific nature, as one of the defining characteristics of this dissertation is the central role that methodological innovations play in every chapter. One could ask whether such a methodological orientation is a good thing. In the end, methodology is just a tool and not an aim in itself. I think that such a methodological dissertation has its place within a substantive field like social stratification research, but such a study should meet a number of challenges. When proposing new methodologies, it is easy to get carried away and to purely focus on applying the latest and most fashionable techniques. Similarly, when pointing out a defect in a methodology it is very easy to forget that all models are defective, as models by their very nature are simplifications of reality and a simplification is nothing other than a 'reasonable error'. In other words it is not enough to show that one can invent or apply new methodologies or show that some 'old' methodology is defective, one must also show that this helps to either better answer existing questions or answer new questions. Moreover, when one proposes new methodologies it is easy to forget that the aim is to create a new tool that can be used by others. If it takes more than a reasonable amount of effort for other researchers to use this new method, then the methodological study has not achieved its aim. In this dissertation I have attempted to meet these challenges by focussing in each chapter on using the methodological innovations to answer substantive questions, leading to some truly new findings, thus showing that it is not just new technology but that this new technology contributes to the study of educational inequality. Moreover, the methods proposed in this dissertation used either existing software or new software was written to implement the new methodologies. In particular, chapter 4 used the `locfit` module by Loader (1999), while two new software modules were written within the statistical programme Stata (StataCorp, 2007) to implement the remaining new methodologies: `seqlogit` (Technical materials II) for chapters 6 and 7 and `propcnsreg` (Technical materials I) for chapters 3 and 5. This has enabled the new methods proposed in this dissertation to be accessible to other users.

Technical materials

Technical materials I

sheafcoef and propcnsreg:

Stata modules for fitting a measurement model with causal indicators

Both chapters 3 and 5 used the `propcnsreg` package, but for subtly different purposes. In chapter 3 information from several educational category dummy variables were combined into a single optimally-scaled education variable, while in chapter 5 tested whether the relative sizes of the effect of several parental background variables have remained constant over time. The aim of this appendix is to describe both this package and a related package: `sheafcoef` (Buis, 2009b). Both `propcnsreg` and `sheafcoef` have been implemented in Stata (StataCorp, 2007).

The models implemented in both packages can be derived from the assumption that the observed variables influence the latent variable. A common alternative assumption is that the latent variable influences the observed variables. For example, factor analysis is based on this alternative assumption. To distinguish between these two situations, some authors, following Bollen (1984) and Bollen and Lennox (1991), call the observed variables "effect indicators" when they are influenced by the latent variable, and they call the observed variables "causal indicators" when they influence the latent variable. Distinguishing between these two is important as each requires a very different strategy for recovering the latent variable and its effect. In a basic (exploratory) factor analysis, which is a model for effect indicators, one assumes that the only thing the indicators have in common is the latent variable, so any correlation between these variables must be due to the latent variable, and it is this correlation that is used to recover the latent variable. In `propcnsreg` and `sheafcoef`, which estimate models for causal indicators, the latent variable is assumed to be a weighted sum of the indicators (and optionally an error term), and the weights are estimated such that they are optimal for predicting the dependent variable. Within the models implemented in the `propcnsreg` package this turns out to be equivalent to a proportionality constraint, that is, the constraint that the relative influence of each indicator remains constant over a set of other variables, in case of Chapter 5, cohort and gender.

Models for dealing with causal indicators come in roughly three flavors: A model with "sheaf coefficients" (Heise, 1972), a model with "parametricaly weighted covariates” (Yamaguchi, 2002), and a Multiple Indicators and Multiple Causes (MIMIC) model (Hauser and Goldberger, 1971). The latter two can be estimated using `propcnsreg`, while the former can be estimated using `sheafcoef`.

I.1 Sheaf coefficient

The sheaf coefficient is the simplest model of the three. Assume we want to explain a variable y using three observed variables x_1, x_2, and x_3, and we think that x_1 and x_2 actually influence y through a latent variable η and x_3 is a control variable. Because η is a latent variable, we need to fix its origin and its unit. The origin can be fixed by setting η to 0 when both x_1 and x_2 are 0, and the unit can be fixed by setting the standard deviation of η equal to 1. The model starts with a multiple regression model, where the βs are the regression coefficients and ε is a normally distributed error term, with a mean of 0 and a standard deviation that is to be estimated:

$$y = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \beta_3 x_3 + \varepsilon \tag{I.1}$$

We want to turn this into equations (I.2) and (I.3), where λ is the effect of the latent variable and the γs are the effects of the observed variables on the latent variable:

$$y = \beta_0 + \lambda\eta + \beta_3 x_3 + \varepsilon \tag{I.2}$$

$$\eta = \gamma_0 + \gamma_1 x_1 + \gamma_2 x_2 \tag{I.3}$$

We can fix the origin of η by constraining γ_0 to be 0. This way η will be 0 when both x_1 and x_2 equal 0. This leaves γ_1 and γ_2. We want to choose values for these parameters such that η optimally predicts y, and the standard deviation of η equals 1. This means that γ_1 and γ_2 are going to be a transformation of β_1 and β_2. We can start with an initial guess that γ_1 equals β_1 and γ_2 equals β_2, and call the resulting latent variable η'. This will get us closer to where we want to be, as we now have values for all parameters: γ_0=0, γ_1'=β_1, γ_2'=β_2, and λ'=1. The value for λ' is derived from the fact that that is the only value where equations (I.2) and (I.3) lead to equation (I.1). However, the standard deviation of η' will generally not be equal to 1. The standard deviation of η' can be calculated as follows:

$$\text{sd}(\eta') = \sqrt{\beta_1^2 \text{var}(x_1) + \beta_2^2 \text{var}(x_2) + 2\beta_1\beta_2 \text{cov}(x_1, x_2)}$$

We can recover η by dividing η' by its standard deviation, which means that the true values of γ_1 and γ_2 are actually $\beta_1/\text{sd}(\eta')$ and $\beta_2/\text{sd}(\eta')$. If we divide η' by its standard deviation, then we must multiply λ' by that same number to ensure that equations (I.2) and (I.3) continue to lead to equation (I.1). As a consequence λ will equal $\text{sd}(\eta')$.

This illustrates how the following set of assumptions can be used to recover the latent variable and its effect on the dependent variable:

- the latent variable is a weighted sum of the observed variables such that the latent variable optimally predicts the dependent variable.
- a constraint that fixes the origin of the latent variable.
- a constraint that fixes the unit of the latent variable.

One possible application of the sheaf coefficient is the comparison of effect sizes of different blocks of variables. For example, we may have a block of variables representing the family situation of the respondent and another block of variables representing characteristics of the work situation and we ask ourselves whether the work situation or the family situation is more important for determining a certain outcome variable. In that case we would estimate a model with two latent variables, one for the family situation and one for the work situation, and since both latent variables are standardized their effects will be comparable.

This can be useful, but a sheaf coefficient merely reorders the information obtained from a regular regression. As a consequence, it is simply a different way of looking at the regression results, and it does not impose a testable constraint. Moreover, this model does not allow for any errors in the measurement of η, as equation (I.3) does not contain an error term.

I.2 Parametricaly weighted covariates

The model with parametricaly weighted covariates Yamaguchi (2002) builds on the model with sheaf coefficients, but allows the effect of the latent variable to change over one or more other variables. This means that equation (I.4), where the effect of η changes over x_3 will be estimated, instead of equation (I.2).

$$y = \beta_0 + (\lambda_0 + \lambda_1 x_3)\eta + \beta_3 x_3 + \varepsilon \tag{I.4}$$

If η is replaced by equation (I.3), and the origin of η is fixed by constraining γ_0 to be zero, we get:

$$\begin{aligned} y &= \beta_0 + (\lambda_0 + \lambda_1 x_3)(\gamma_1 x_1 + \gamma_2 x_2) + \beta_3 x_3 + \varepsilon \\ &= \beta_0 + (\lambda_0 + \lambda_1 x_3)\gamma_1 x_1 + (\lambda_0 + \lambda_1 x_3)\gamma_2 x_2 + \beta_3 x_3 + \varepsilon \end{aligned}$$

This means the effect of x_1 (through η) on y equals $(\lambda_0 + \lambda_1 x_3)\gamma_1$, and that the effect of x_2 (through η) on y equals $(\lambda_0 + \lambda_1 x_3)\gamma_3$. This implies the following constraint: for every value of x_3, the effect of x_1 relative to x_2 will always be $(\lambda_0 + \lambda_1 x_3)\gamma_1/(\lambda_0 + \lambda_1 x_3)\gamma_2 = \gamma_1/\gamma_2$, which is a constant. In other words, the model with parametricaly weighted covariates imposes a proportionality constraint.

This proportionality constraint can also be of substantive interest without referring to a latent variable. Consider a model where one wants to explain the respondent's education (ed) with the eduction of the father (fed) and the mother (med), and that one is interested in testing whether the relative contribution of the mother's education has increased over time. `propcnsreg` will estimate this model under the null hypothesis that the relative contributions of fed and med have remained constant over time. Notice that the effects of fed and med are allowed to change over time, but the effects of fed and med are constrained to change by the same proportion over time. So if the effect of fed drops by 10% over a decade, then so does the effect of med.

The default way in which `propcnsreg` will identify the unit of the latent variable is by setting its standard deviation to 1. Alternatively, the unit can be identified in one of the following two ways: the coefficient λ_0 can be set to 1, which means that γ_1 and γ_2 represent the indirect effects of x_1 and x_2 through the latent variable on y when x_3 equals 0. This is the default parametrization, but can also be explicitly requested by specifying the `lcons` option. Alternatively, either the coefficient γ_1 or γ_2 can be set to 1, which means that the unit of the latent variable will equal the unit of x_1 or x_2 respectively. This can be done by specifying the `unit(`*varname*`)` option.

I.3 MIMIC

The MIMIC model builds on the model with parametricaly weighted covariates by assuming that the latent variable is measured with error. This means that the following model is estimated:

$$y = \beta_0 + (\lambda_0 + \lambda_1 x_3)\eta + \beta_3 x_3 + \varepsilon_y \tag{I.5}$$

$$\eta = \gamma_1 x_1 + \gamma_2 x_2 + \varepsilon_\eta \tag{I.6}$$

Where ε_y and ε_η are independent normally distributed error terms with means zero and standard deviations that need to be estimated. By replacing η in equation (I.5) by equation (I.6) one can see that the error term of this model is:

$$\varepsilon_y + (\lambda_0 + \lambda_1 x_3)\varepsilon_\eta$$

This combined error term will also be normally distributed, as the sum of two independent normally distributed variables is itself also normally distributed. The mean of this combined error term will be zero and it will have the following standard deviation:

$$\sqrt{\text{var}(\varepsilon_y) + (\lambda_0 + \lambda_1 x_3)^2 \text{var}(\varepsilon_\eta)}$$

So the empirical information that is used to separate the standard deviation of ε_y from the standard deviation of ε_η, is the changes in the residual variance over x_3. The data will thus contain rather indirect information that can be used for estimating this model. However, if the model is correct, it will make it possible to control for measurement error in the latent variable.

There is an important downside to this model, and that is that heteroscedasticity, and in particular changes in the variance of ε_y over x_3, could have a distorting influence on the parameter estimates of λ_0 and λ_1. Consider again the example of wanting to explain the respondent's education through the education of the father and the mother, but now assume that we are interested in how the effect of the latent parental education variable changes over time. In this case we have good reason to suspect that the variance of ε_y will also change over time: education consists of a discrete number of categories, and in early cohorts most of the respondents tend to cluster in the lowest categories. Over time the average level of education tends to increase, which in practice means that the respondents tend to cluster less in the lowest category, and have more room to differ from one another. As a consequence the residual variance is likely to have increased over cohorts. Normally this heteroscedasticity would not be an issue of great concern, but in a MIMIC model this heteroscedasticity is incorrectly interpreted as indicating that there is measurement error in the latent variable representing parental education. Moreover, this 'information' on the measurement error is used to 'refine' the estimates of λ_0 and λ_1. So, this would be an example where the MIMIC model would not be appropriate.

I.4 Maximization of the likelihood function

A difficulty with both the model with parametricaly weighted covariates and the MIMIC model is that the parameters are highly correlated, thus making it difficult for the standard maximization algorithms to find the maximum of the likelihood function. To overcome this issue, an EM algorithm is first used to find suitable starting values. The EM algorithm breaks the correlation by first treating the weights for the observed variables as fixed and estimating the effect of the latent variable, and then treating the effect of the latent variable as fixed and estimating the weights. By default, this is iterated 20 times or until convergence. These parameter estimates are then used as starting values for the regular maximum likelihood algorithm.

I.5 Example

The `sheafcoef` programme uses the fact that a sheaf coefficient is simply a transformation of regression coefficients, which allows it to be implemented as a post-estimation programme. This means that one must first estimate a regression model, using an estimation command like `regress` or `logit`, and then one can use `sheafcoef` to redisplay the results as a model with sheaf coefficients. It is therefore possible to use `sheafcoef` for continuous, ordered, and binomial dependent variables.

The use of this command can be illustrated using the `nlsw88.dta` dataset that comes with Stata (StataCorp, 2007). The first step is to open that dataset using the `sysuse` command, and prepare the variables.

```
. sysuse nlsw88, clear
(NLSW, 1988 extract)
.
. gen highschool = grade == 12 if grade < .
(2 missing values generated)
. gen somecollege = grade > 12 & grade < 16 if grade < .
(2 missing values generated)
. gen college = grade >= 16 if grade < .
(2 missing values generated)
.
. gen lnwage = ln(wage)
.
. gen ttl_exp2 = ttl_exp^2
.
. gen white = race == 1 if race < .
. gen other = race == 3 if race < .
```

In this example we have a set of dummies representing an individuals education (*highschool*, *somecollege*, and *college*, meaning that the reference category is those that have not finished high school), and a set of dummies representing an individual's race (*white*, and *other*, with African Americans as reference category), and we wonder which set of variables is more important for predicting an individuals wage while controlling for total experience in the labor market (*ttl_exp* and *ttl_exp2*). So, first a regression of all these variables on log wage is estimated. After that, `sheafcoef` is used, specifying in the `latent()` option the blocks of variables that belong to the same latent variable. The blocks are separated using a semi-colon (`;`). Each block of variables is preceded by its name followed by a colon (`:`). So in this example, the block of education dummies is given the name `educ` and the block of race dummies is given the name `race`. The parameters `educ` and `race` in the `main` equation represent the effects of the two latent variables. The parameters in the `on_educ` and `on_race` equations represent the effects of the dummies on the education and race latent variable respectively. The results show that education is more important than race for determining a person's income.

```
. reg lnwage white other ttl_exp ttl_exp2 highschool somecollege college

      Source |       SS       df       MS              Number of obs =    2244
-------------+------------------------------           F(  7,  2236) =  120.98
       Model |  203.545105     7  29.0778722           Prob > F      =  0.0000
    Residual |  537.417694  2236  .240347806           R-squared     =  0.2747
-------------+------------------------------           Adj R-squared =  0.2724
       Total |  740.962799  2243  .330344538           Root MSE      =  .49025

------------------------------------------------------------------------------
      lnwage |      Coef.   Std. Err.      t    P>|t|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
       white |    .118158   .0241501     4.89   0.000     .0707991    .1655169
       other |   .1079395   .0987438     1.09   0.274    -.0856996    .3015786
     ttl_exp |   .0616495    .009803     6.29   0.000     .0424255    .0808734
    ttl_exp2 |  -.0008656    .000395    -2.19   0.029    -.0016403   -.0000909
  highschool |   .1087398   .0320975     3.39   0.001     .0457958    .1716838
 somecollege |   .3568001   .0365223     9.77   0.000     .2851789    .4284213
     college |   .5167365   .0360893    14.32   0.000     .4459644    .5875086
       _cons |   .9272152    .061262    15.14   0.000      .807079    1.047351
------------------------------------------------------------------------------

. sheafcoef, latent(educ: highschool somecollege college ; race: white other)
------------------------------------------------------------------------------
      lnwage |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
main         |
        educ |   .1898757   .0107139    17.72   0.000     .1688769    .2108746
        race |   .0517396   .0105663     4.90   0.000       .03103    .0724491
     ttl_exp |   .0616495    .009803     6.29   0.000     .0424359     .080863
    ttl_exp2 |  -.0008656    .000395    -2.19   0.028    -.0016399   -.0000913
       _cons |   .9272152    .061262    15.14   0.000      .807144    1.047286
-------------+----------------------------------------------------------------
on_educ      |
  highschool |   .5726894   .1645151     3.48   0.000     .2502457    .8951332
 somecollege |   1.879124   .1570053    11.97   0.000       1.5714    2.186849
     college |   2.721446   .1063338    25.59   0.000     2.513036    2.929857
-------------+----------------------------------------------------------------
on_race      |
       white |   2.283707   .0197364   115.71   0.000     2.245024     2.32239
       other |   2.086208   1.859547     1.12   0.262    -1.558437    5.730853
------------------------------------------------------------------------------
```

The `propcnsreg` programme can estimate both models with parametricaly weighted covariates and MIMIC models. Unlike the models with sheaf coefficients, these models need to be separately estimated, and can thus not be as flexibly implemented as the post-estimation command `sheafcoef`. In particular, `propcnsreg` can only be used for continuous dependent variables with (approximately) normally distributed errors.

The use of `propcnsreg` can be illustrated by continuing the example. Now we assume that the effect of education changes for different levels of experience. The parameters in the 'constrained' panel represent the scale of education, such that parameters of high school and some college represent the positions of these levels relative to less than high school (0) and college (1). These are the effects of the education dummies on the latent variable. The parameters in the panel 'lambda' represent how the effect of the latent optimally-scaled education changes when experience changes. The unconstrained panel shows the main effects of experience and the control variables. A test of the proportionality constraint is reported at the bottom of the output.

```
. propcnsreg lnwage white other ttl_exp ttl_exp2, /*
>         */ lambda(ttl_exp ttl_exp2) /*
>         */ constrained(highschool somecollege college) /*
>         */ unit(college) nolog

                                                  Number of obs   =       2244
                                                  LR chi2(10)     =     101.57
Log likelihood = -1573.1308                       Prob > chi2     =     0.0000


Constraint: [constrained]college = 1
------------------------------------------------------------------------------
      lnwage |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
unconstrai~d |
       white |   .1166583   .0240475     4.85   0.000     .0695259    .1637906
       other |   .1101377   .0981958     1.12   0.262    -.0823226    .3025981
     ttl_exp |   .0211701    .015895     1.33   0.183    -.0099836    .0523237
    ttl_exp2 |   .0006399   .0006602     0.97   0.332    -.0006541    .0019339
       _cons |   1.150775   .0859399    13.39   0.000     .9823357    1.319214
-------------+----------------------------------------------------------------
constrained  |
  highschool |   .2431708   .0550686     4.42   0.000     .1352384    .3511032
 somecollege |   .7056825   .0538163    13.11   0.000     .6002046    .8111605
     college |          1          .        .       .            .           .
-------------+----------------------------------------------------------------
lambda       |
     ttl_exp |   .1079688   .0299633     3.60   0.000     .0492419    .1666957
    ttl_exp2 |  -.0039162   .0012162    -3.22   0.001    -.0062999   -.0015325
       _cons |  -.1390864   .1748364    -0.80   0.426    -.4817595    .2035867
-------------+----------------------------------------------------------------
ln_sigma     |
       _cons |  -.7178998    .014927   -48.09   0.000    -.7471563   -.6886434
------------------------------------------------------------------------------
LR test vs. unconstrained model: chi2(4) =     13.31   Prob > chi2 =     0.010
```

A MIMIC model can be estimated using `propcnsreg` by specifying the `mimic` option. This means that an extra parameter (`ln_sigma_latent`) is estimated representing the log of the standard deviation of the measurement error of the latent variable. In this case this does not lead to major changes in the results.

```
. propcnsreg lnwage white other ttl_exp ttl_exp2, /*
>         */ lambda(ttl_exp ttl_exp2) /*
>         */ constrained(highschool somecollege college) /*
>         */ unit(college) mimic nolog

                                                  Number of obs   =       2244
                                                  LR chi2(10)     =     135.26
Log likelihood = -1571.1459                       Prob > chi2     =     0.0000


Constraint: [constrained]college = 1
------------------------------------------------------------------------------
      lnwage |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
unconstrai~d |
       white |   .1187835   .0240303     4.94   0.000     .0716851    .1658819
       other |   .1017959   .0970624     1.05   0.294    -.0884429    .2920346
     ttl_exp |    .016967   .0151072     1.12   0.261    -.0126427    .0465766
    ttl_exp2 |   .0009142   .0006198     1.48   0.140    -.0003006    .0021291
       _cons |   1.160823   .0805208    14.42   0.000     1.003005    1.318641
-------------+----------------------------------------------------------------
constrained  |
  highschool |   .2304431   .0556791     4.14   0.000     .1213141    .3395721
 somecollege |   .7022772   .0550723    12.75   0.000     .5943375    .8102169
     college |          1          .        .       .            .           .
-------------+----------------------------------------------------------------
lambda       |
     ttl_exp |   .1182191   .0291939     4.05   0.000     .0610002    .1754381
    ttl_exp2 |    -.00456   .0011831    -3.85   0.000    -.0068787   -.0022412
       _cons |   -.160952    .168419    -0.96   0.339    -.4910472    .1691433
-------------+----------------------------------------------------------------
ln_sigma     |
       _cons |  -.8121788   .0501504   -16.19   0.000    -.9104719   -.7138858
-------------+----------------------------------------------------------------
ln_sigma_l~t |
       _cons |  -.9502158   .2506895    -3.79   0.000    -1.441558   -.4588735
------------------------------------------------------------------------------
```

This example illustrates how to estimate models with the three types of causal indicators using the `sheafcoef` and `propcnsreg` modules in Stata. A complete description of the syntax of `sheafcoef` and `seqlogit` is given below.

I.6 Syntax and options

Syntax of `sheafcoef`

`sheafcoef, latent(`*varlist_1* `[ ;` *varlist_2* `[;` *varlist_3* `[...]]] )`
`[ eform post iterate(#) level(#) ]`

Options of `sheafcoef`

`latent(`*varlist_1* `[;`*varlist_2* `[;` *varlist_3* `[...]]])` specifies the blocks of variables that make up the latent variables, whereby each block is separated by a semicolon (;). Each block needs to consist of at least two variables. These variables must be explanatory variables in the estimation command preceding `sheafcoef`, and the same variable can only appear in one block.

`eform` specifies that the effects of the latent variable and the control variables are exponentiated. The effects of the indicator variables in each block on its latent variable are not exponentiated, because these represent the effects of these variables on the standardized latent variable and not on the dependent variable. This option can be useful after commands like `logit` or `poisson`, as this will cause the effects on the dependent variables to be displayed in the form of odds ratios and incidence rate ratios respectively.

`post` causes `sheafcoef` to behave like a Stata estimation (e-class) command. When `post` is specified, `sheafcoef` will post the vector of transformed estimators and its estimated variance-covariance matrix to `e()`. This option, in essence, makes the transformation permanent. Thus you could, after posting, treat the transformed estimation results in the same way as you would treat results from other Stata estimation commands. For example, after posting, you could use `test` to perform simultaneous tests of hypotheses on linear combinations of the transformed estimators.

Specifying `post` clears the previous estimation results, which can then only be recovered by refitting the original model or by storing the estimation results before running `sheafcoef` and then restoring them; see [R] **estimates store**[1].

`level(#)` specifies the confidence level, as a percentage, for confidence intervals. The default is `level(95)` or as set by `set level`, see [R] **level**.

`iterate(#)` specifies the maximum number of iterations used to find the optimal step size in calculating numerical derivatives of the transformations with respect

[1] I am following Stata's convention when referencing to the manuals of Stata. These conventions are discussed in the User's Guide that comes with Stata, section 1.2.2: [U] **1.2.2 Cross-referencing**.

to the original parameters. By default, the maximum number of iterations is 100, but convergence is usually achieved after only a few iterations. You should rarely have to use this option.

Syntax of `propcnsreg`

`propcnsreg` *depvar* [*indepvars*] [*if*] [*in*] [*weight*] `,` `constrained(`*varlist*`)` `lambda(`*varlist*`)` [`standardized` `lcons` `unit(`*varname*`)` `mimic` `robust` `cluster(`*varname*`)` `level(`*#*`)` *em_maximize_options* *maximize_options*]

Options of `propcnsreg`

`constrained(`*varlist_c*`)` specifies the variables that are measurements of the same latent variable. The effects of these variables are to be constrained to change by the same proportion as the variables specified in `lambda()` change.

`lambda(`*varlist_l*`)` specifies the variables along which the effects of the latent variable changes.

`standardized` specifies that the unit of the latent variable is identified by constraining the standard deviation of the latent variable to be equal to 1. This is the default parametrization.

`lcons` specifies that the parameters of the variables specified in the option `constrained()` measure the indirect effect of these variables through the latent variable on the dependent variable when all variables specified in the option `lambda()` are zero.

`unit(`*varname*`)` specifies that the scale of the latent variable is identified by constraining the unit of the latent variable to be equal to the unit of *varname*. The variable *varname* must be specified in `constrained()` option.

`mimic` specifies that a MIMIC model is to be estimated.

`robust` specifies that the Huber/White/sandwich estimator of variance is to be used instead of the traditional calculation; see [U] **23.14 Obtaining robust variance estimates**. `robust` combined with `cluster()` allows observations which are not independent within cluster (although they must be independent between clusters).

`cluster(`*clustervar*`)` specifies that the observations are independent across groups (clusters) but not necessarily within groups. *clustervar* specifies to which group each observation belongs; e.g., `cluster(`*personid*`)` in data with repeated observations on individuals. See [U] **23.14 Obtaining robust variance estimates**. Specifying `cluster()` implies `robust`.

`level(#)` specifies the confidence level, in percent, for the confidence intervals of the coefficients; see [R] **level**.

em_maximize_options

`emiterate(#)` specifies the maximum number of iterations for the EM algorithm. When the number of iterations equals `emiterate()`, the EM algorithm stops. If convergence is declared before this threshold is reached, it will stop when convergence is declared. The default value of `emiterate()` is 20.

`emtolerance(#)` specifies the tolerance for the coefficient vector. When the relative change in the coefficient vector from one iteration to the next is less than or equal to `emtolerance()`, the `emtolerance()` convergence criterion is satisfied. `emtolerance(1e-6)` is the default.

`emltolerance(#)` specifies the tolerance for the log likelihood. When the relative change in the log likelihood from one iteration to the next is less than or equal to `emltolerance()`, the `emltolerance()` convergence is satisfied. `emltolerance(1e-7)` is the default

These options are seldom used.

maximize_options

`difficult`, `technique(`*algorithm_spec*`)`, `iterate(#)`, `trace`, `gradient`, `showstep`, `hessian`, `shownrtolerance`, `tolerance(#)`, `ltolerance(#)`, `gtolerance(#)`, `nrtolerance(#)`, `nonrtolerance(#)`; see [R] **maximize**. These options are seldom used.

Technical materials II

---

seqlogit:
Stata module for fitting a sequential logit model

II.1 Introduction

Chapters 6 and 7 propose two extensions to the sequential logit model, both of which have been implemented in Stata (StataCorp, 2007) as the `seqlogit` package (Buis, 2007b). The aim of this appendix is to show how to use this package. This will be done by presenting an example analysis using data already present in Stata and by giving a complete description of its syntax.

II.2 Example

The use of the `seqlogit` package is illustrated using the `nlsw88.dta` dataset that comes with Stata, and can be opened using the `sysuse` command. This dataset contains a variable *grade* measuring the respondent's highest achieved level of education in years. The dependent variable is created by transforming the variable *grade* into the variable *ed*, which measures the respondent's highest achieved level of education in the categories: less than high school (1), high school (2), some college (3), and college (4). In the example I assume that the respondents achieved their level of education by passing or failing the following sequence of transitions:

1. respondents either finished high school or not
2. those respondents that finished high school either went to college or not
3. those respondents that went to college either finished a four-year course or not

This decision tree is fed into `seqlogit` using the `tree()` option. Within this option the levels are represented with the values in the dependent variable (*ed*), the transitions are separated using commas, and the choices within a transition are separated by a colon. This tree would thus be represented in the following way: `tree(1 : 2 3 4 , 2 : 3 4 , 3 : 4)`. So the first transition is a choice between less than high school (1) and all other levels (2, 3, and 4), the second transition is a choice between leaving after high school (2) versus going to college (3 and 4), and the final transition is a choice between some college (3) and a four-year course (4)[1].

The key explanatory variable in this example is whether or not the respondent is white (*white*), and the effect of this variable can change over time (*byr*). These variables need to be specified in the `offinterest()` and `over()` options respectively in order to make use of the post-estimation commands that come with `seqlogit`. This will cause `seqlogit` to make a new variable *_white_X_byr*, the interaction term between *white* and *byr*, and to add the variables *white* and *_white_X_byr* to the list of explanatory variables. Notice that the main effect of *byr* is not added automatically and needs to be added separately as one of the independent variables. This makes it possible for the main effect of *byr* to have a different functional form than the interaction effect. The values assigned to each level of education are specified in the `levels` option. This won't influence the output obtained from `seqlogit`, but will influence post-estimation commands like `predict` and `seqlogitdecomp`. Finally, I added a variable indicating whether or not the respondent lived in the south of the USA (*south*) as a control variable, and I added the `or` option to specify that the odds ratios are to be displayed. Together this resulted in the following sequence of commands and output:

```
. sysuse nlsw88, clear
(NLSW, 1988 extract)

. gen ed = cond(grade< 12, 1, ///
>          cond(grade==12, 2, ///
>          cond(grade<16,3,4))) if grade < .
(2 missing values generated)

. gen byr = (1988-age-1950)/10

. gen white = race == 1 if race < .
```

(*Continued on next page*)

[1]The choices specified in the `tree()` option do not have to be binary (pass or fail). For example, we may believe that after finishing high school, students choose between leaving the schooling system, junior college, and college. In that case the `tree()` option would look like `tree(1 : 2 3 4, 2 : 3 : 4)`.

```
. seqlogit ed byr south,                      ///
>          ofinterest(white) over(byr)        ///
>          tree(1 : 2 3 4, 2 : 3 4, 3 : 4) ///
>          levels(1=6, 2=12, 3=14, 4= 16)  ///
>          or nolog

Transition tree:

Transition 1: 1 : 2 3 4
Transition 2: 2 : 3 4
Transition 3: 3 : 4

Computing starting values for:

Transition 1
Transition 2
Transition 3

                                                  Number of obs   =       2244
                                                  LR chi2(12)     =     110.38
Log likelihood = -2881.2013                       Prob > chi2     =     0.0000

------------------------------------------------------------------------------
          ed | Odds Ratio   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
_2_3_4v1     |
         byr |   3.377124   1.062584     3.87   0.000     1.822735    6.257061
       south |   .6440004   .0807557    -3.51   0.000     .5036723    .8234254
       white |    2.17841   .3029219     5.60   0.000     1.658726    2.860913
_white_X_byr |   .3330505   .1351488    -2.71   0.007     .1503489    .7377681
-------------+----------------------------------------------------------------
_3_4v2       |
         byr |   1.139388   .3722391     0.40   0.690     .6005969    2.161523
       south |   .8258418   .0793651    -1.99   0.046     .6840607     .997009
       white |   1.090765   .1244936     0.76   0.447     .8721274    1.364214
_white_X_byr |   .9148277   .3372712    -0.24   0.809     .4441455    1.884314
-------------+----------------------------------------------------------------
_4v3         |
         byr |   1.217693   .5757529     0.42   0.677     .4820255    3.076134
       south |   1.501026   .2063442     2.95   0.003     1.146501    1.965178
       white |   1.340784   .2183215     1.80   0.072     .9744438     1.84485
_white_X_byr |   .7585029   .4037806    -0.52   0.604     .2671958    2.153203
------------------------------------------------------------------------------
```

The results show that being white was particularly beneficial at the first transition (whether or not to finish high school), but had little effect at the higher transition. The effect of being white decreased only at the first transition. Chapter 6 showed that one can also derive the effect on the highest achieved level of education from this sequential logit model if we can assign a value to each level of education. Within Stata, this effect can be recovered after `seqlogit` using the `predict` command with the `effect` option. In this example, the levels are given values that approximately equal the years of education. One characteristic of this effect is that it will change when any of the explanatory variables change, so in order to show how the effect of `white` changed over time we first need to create a dataset where people only differ with respect to time. This is done in the example below by setting *south* and *white* to 0. The `preserve` and `restore` commands are used to ensure that these changes are only temporary. This sequence of code results in Figure II.1, which shows that the advantage of being white dropped from almost 3 years to about .5 years.

```
. preserve

. replace white = 0
(1637 real changes made)

. replace _white_X_byr = white * byr
(1479 real changes made)

. replace south = 0
(942 real changes made)

.
. predict eff, effect

.
. gen coh = byr *10 + 1950

. label variable coh "year of birth"

.
. twoway line eff coh, sort                          ///
>        ylab(0(1)3)                                 ///
>        ytitle("effect of respondent being white")

. restore
```

Figure II.1: Effect of the respondents being white on their highest achieved level of education

Chapter 6 showed that this effect on the highest achieved level of education is a weighted sum of the effects on passing each transition. The contribution of each transition can thus be visualized by the area of a rectangle with a width equal to the weight and a height equal to the effect on the probability of passing the transition (the log odds ratio). This is shown in Figure II.2 for three different cohorts, showing that the contribution of the first transition to the effect on the highest attained level of education has dropped dramatically over time. This graph was made using the call to `seqlogitdecomp` command shown below. The `at()` option tells that the effect is being decomposed for black respondents who are not from the south. The `overat()` option tells that this decomposition is shown for the cohorts -.5, 0, and .4. Time is measured in decades since 1950, and the `subtitle()` option is used to give more meaningful column titles. The transitions are labelled using the `eqlabel()` option. Each transition label spans two lines. This is achieved by surrounding each line with double quotes (`" "`). Each transition's label is in turn surrounded by so-called compound quotes (`` `" "' ``), to tell Stata which lines belong together.

```
. seqlogitdecomp, at(south 0 white 0)                ///
>          overat(byr -.5, byr 0, byr .4)            ///
>          subtitle("1945" "1950" "1954")            ///
>          eqlabel(‘""finish" "high school""’        ///
>                  ‘""high school v" "some college""’ ///
>                  ‘""some college v" "college""’)    ///
>          xline(0) yline(0)
```

Chapter 6 also showed that the weights can in turn be decomposed as the product of three elements: The proportion of respondents at risk, the variance of the dummy variable indicating whether one passes a transition or not, and the expected gain in level of education resulting from passing. The weights and each of these elements can also be recovered using `predict` and can be displayed using the same tricks as were used when displaying the effect of `white` on the highest achieved level of education.

Chapter 7 proposed a way of assessing how sensitive the results of a sequential logit model is to unobserved heterogeneity. This strategy consists of estimating the effects of the observed variables given various scenarios concerning the degree of unobserved heterogeneity. The unobserved heterogeneity is assumed to be the result of a normally distributed unobserved variable, and the degree of unobserved heterogeneity is captured by the standard deviation of this variable. This is implemented in the `seqlogit` package in the form of the `sd()` option, which sets the standard deviation of the unobserved variable. The entire sensitivity analysis consists of multiple models, each with a different degree of unobserved heterogeneity. In the example below, just one such model is shown. Notice that we first need to drop the *_white_X_byr* variable, as this variable will be created by each call to `seqlogit`, and this will cause an error if the variable already exists.

Figure II.2: Decomposition of effect of the respondents being white on their highest achieved level of education

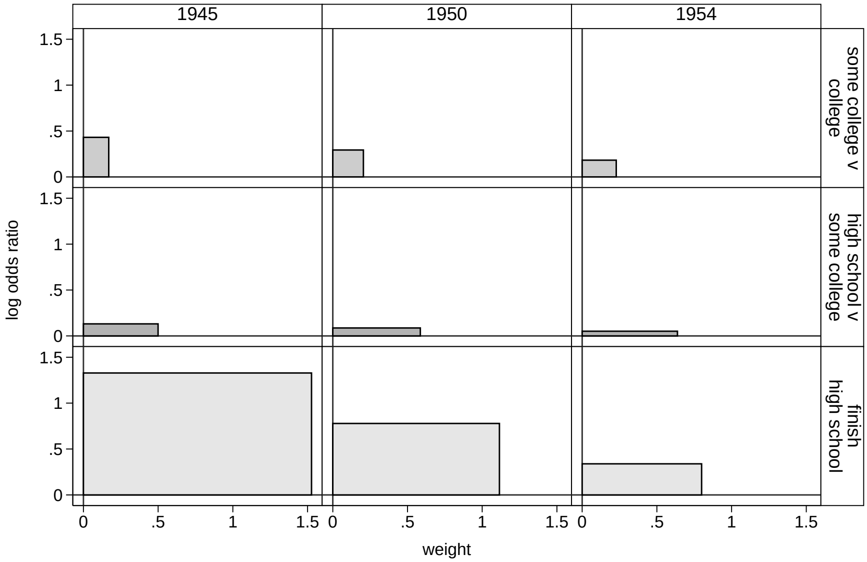


```
. drop _white_X_byr

. seqlogit ed byr south,                   ///
>          ofinterest(white) over(byr)     ///
>          tree(1 : 2 3 4, 2 : 3 4, 3 : 4) ///
>          or sd(1) nolog

Transition tree:

Transition 1: 1 : 2 3 4
Transition 2: 2 : 3 4
Transition 3: 3 : 4

Computing starting values for:

Transition 1
Transition 2
Transition 3
```

(*Continued on next page*)

```
                                                  Number of obs   =       2244
                                                  LR chi2(12)     =     109.61
Log likelihood = -2881.5827                       Prob > chi2     =     0.0000

------------------------------------------------------------------------------
          ed | Odds Ratio   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
_2_3_4v1     |
         byr |   4.143983   1.527877     3.86   0.000     2.011787    8.535991
       south |   .6132818   .0872595    -3.44   0.001     .4640328    .8105343
       white |   2.429149   .3848478     5.60   0.000     1.780734    3.313669
_white_X_byr |   .2739024   .1274384    -2.78   0.005     .1100418    .6817639
-------------+----------------------------------------------------------------
_3_4v2       |
         byr |   1.387615   .5412719     0.84   0.401     .6460077    2.980576
       south |   .7622253   .0879274    -2.35   0.019     .6079838    .9555969
       white |   1.212081   .1658456     1.41   0.160     .9269668     1.58489
_white_X_byr |   .7634598   .3365498    -0.61   0.540     .3217791      1.8114
-------------+----------------------------------------------------------------
_4v3         |
         byr |   1.481788     .82091     0.71   0.478     .5002889    4.388857
       south |    1.49327   .2411018     2.48   0.013     1.088189    2.049144
       white |   1.539153   .2944113     2.25   0.024     1.057944    2.239241
_white_X_byr |   .6233829   .3892371    -0.76   0.449      .183345    2.119536
------------------------------------------------------------------------------
The effect of the standardized unobserved variable is fixed at:
--------------------
equation    | sd
------------+-------
_2_3_4v1    | 1
_3_4v2      | 1
_4v3        | 1
--------------------
```

Part of the problem with unobserved variables is that the distribution of that variable tends to change over transitions due to selection. This change in distribution can be shown using the `uhdesc` command. Of particular interest is the change in the correlation between the observed variable of interest (specified in the `ofinterest()` option in `seqlogit`) and the unobserved variable. This correlation is labelled as `corr(e, x))` in the output. It shows that over transitions an initially non-confounding variable has become a confounding variable. The syntax is similar to the `seqlogitdecomp` command.

(*Continued on next page*)

```
. uhdesc, at(south 0 white 0)                 ///
>         overat(byr -.5, byr 0, byr .4)      ///
>         overlab("1945" "1950" "1954")

             | p(atrisk)    mean(e)      sd(e)  corr(e,x)
-------------+-------------------------------------------
1945         |
     2 3 4v1 |     1.000     -0.000      1.000     -0.000
       3 4v2 |     0.706      0.248      0.928     -0.070
         4v3 |     0.347      0.618      0.862     -0.075
-------------+-------------------------------------------
1950         |
     2 3 4v1 |     1.000     -0.000      1.000     -0.000
       3 4v2 |     0.815      0.161      0.945     -0.035
         4v3 |     0.414      0.530      0.876     -0.039
-------------+-------------------------------------------
1954         |
     2 3 4v1 |     1.000     -0.000      1.000     -0.000
       3 4v2 |     0.879      0.108      0.959     -0.012
         4v3 |     0.461      0.475      0.887     -0.015
```

This example illustrates the use of the `seqlogit` package and its post-estimation commands. The full syntax of these commands is described below.

II.3 Syntax and options

Syntax for `seqlogit`

`seqlogit` *depvar* [*indepvars*] [*if*] [*in*] [*weight*] `, tree(`*tree*`)` [`ofinterest(`*varname*`)` `over(`*varlist*`)` `sd(`*#*`)` `rho(`*#*`)` `draws(`*#*`)` `drawstart(`*#*`)` `or` `levels(`*levellist*`)` `constraints(`*numlist*`)` `robust` `cluster(`*varname*`)` `nolog` `level(`*#*`)` *maximize_options*]

Options for `seqlogit`

`tree(`*tree*`)` specifies the sequence of transitions that make up the model. The transitions are separated by commas and the choices within transitions are separated by colons. The levels are represented by the levels of the *depvar*. It is thus convenient to code *depvar* as a series of integers. For example, say there are three levels, 1, 2, and 3, and the first transition consists of a choice between value 1 versus values 2 and 3, and the second transition consists (for those who didn't choose value 1) of a choice between values 2 and 3. The tree option should then be: `tree(1 : 2 3 , 2 : 3)`.

All values of *depvar* must be specified in the tree and all values in the tree must occur in *depvar*. Furthermore, all levels must be accessible through one and only one path through the tree.

`ofinterest(`*varname*`)` specifies the variable whose effect will be decomposed when using the `seqlogitdecomp` command. The variable specified is added to the list of explanatory variables.

`over(`*varlist*`)` specifies the variable(s) over which the effect of the variable specified in the `ofinterest()` option is allowed to change. This/these variable(s) and the interaction effect between the variable(s) specified in `over()` and `ofinterest()` are added to the list of explanatory variables. `ofinterest()` needs to be specified when specifying `over()`.

`sd(`*#*`)` specifies the initial standard deviation of the unobserved variable. The default is 0, which means that there is no unobserved variable.

`rho(`*#*`)` specifies the initial correlation of the unobserved variable and the variable specified in `ofinterest()`. The default is 0, which means that the unobserved variable is initially not a confounding variable.

`draws(`*#*`)` specifies the number of pseudo random draws per observation used when calculating the simulated likelihood. These pseudo random draws are created using a Halton sequence (see: [M-5] **halton()**[2]). The default is 100. Because maximum simulated likelihood is only used when the `sd()` option is specified, the `draws()` option can only be specified when the `sd()` option is specified.

`drawstart(`*#*`)` specifies the index at which the Halton sequence starts. The default is 15. This option can only be specified in combination with the `sd()` option.

`levels(`*levellist*`)` specifies the values attached to each level of the dependent variable. If it is not specified, the values of the dependent variable will be used. The syntax for *levels* is: *#*= *#*[, *#*= *#*, ...]

`or` report odds ratios

`constraints(`*numlist*`)` specifies linear constraints to be applied during estimation, see [R] **constraint**.

`robust` specifies that the Huber/White/sandwich estimator of variance is to be used instead of the traditional calculation; see [U] **23.14 Obtaining robust variance estimates**. `robust` combined with `cluster()` allows observations which are not independent within clusters (although they must be independent between clusters).

[2] I am following Stata's convention when referencing to the manuals of Stata. These conventions are discussed in the User's Guide that comes with Stata, section 1.2.2: [U] **1.2.2 Cross-referencing**.

`cluster(`*clustervar*`)` specifies that the observations are independent across groups (clusters) but not necessarily within groups. *clustervar* specifies to which group each observation belongs; e.g., `cluster(`*personid*`)` in data with repeated observations on individuals. See [U] **23.14 Obtaining robust variance estimates**. Specifying `cluster()` implies `robust`.

`level(`*#*`)` specifies the confidence level, in percent, for the confidence intervals of the coefficients; see [R] **level**.

`nolog` suppresses an iteration log of the log likelihood

maximize_options

`difficult`, `technique(`*algorithm_spec*`)`, `iterate(`*#*`)`, `trace`, `gradient`, `showstep`, `hessian`, `shownrtolerance`, `tolerance(`*#*`)`, `ltolerance(`*#*`)`, `gtolerance(`*#*`)`, `nrtolerance(`*#*`)`, `nonrtolerance(`*#*`)`; see [R] **maximize**. These options are seldom used.

Syntax for `seqlogitdecomp`

`seqlogitdecomp , overat(`*overatlist*`)` [`at(`*atlist*`)`
`subtitle(`*titlelist*`)` `eqlable(`*labellist*`)` `xline(`*linearg*`)`
`yline(`*linearg*`)` `title(`*title*`))` `name(`*name* `[,` *replace*`])`
`xlabel(`*rule or values*`)` `ylabel(`*rule or values*`)`
`yscale(`*axis_suboptions*`)` `xscale(`*axis_suboptions*`)` `ysize(`*#*`)`
`xsize(`*#*`)`]

Options for `seqlogitdecomp`

Specifying the groups to be compared

`overat(`*overlist*`)` Specifies the values of the explanatory variables of the groups that are to be compared. It overrides any value specified in the `at()` option. Each comparison is separated by a comma. The syntax for *overlist* is:

*varname*_1 #[*varname*_2 #[...]], *varname*_1 #[*varname*_2 #[...]], [...]

`at(`*atlist*`)` specifies the values at which the equations are evaluated. The syntax for *atlist* is: *varname*_1 #[*varname*_2 #...]. The equations will be evaluated at the mean values of any of the variables not specified in `at()` or `overat()`.

Say the dependent variable is highest achieved level of education, which is influenced by child's Socioeconomic Status (*ses*) and cohort (*coh*), and the interaction between *ses* and *coh* (*_ses_X_coh*). We want to compare the decomposition of the effect of *ses* over different cohorts for mean value of *ses*. Say that *coh* has only three values: 1, 2, and 3 and the mean value of *ses* is .5. Then the `overat()` and `at()` options would read:

```
overat( coh 1, coh 2, coh 3 ) at( ses .5 )
```

Notice that the values for the interaction term need not be specified in the `overat()` option, as long as it was created using the `over()` option in `seqlogit`.

Other options

`subtitle(`*titlelist*`)` specifies the titles above each group, cohort in the example above. The syntax of *titlelist* is "string" "string" [...]. The number of titles must equal the number of groups.

`eqlabel(`*labellist*`)` specifies labels for each transition. The syntax of *labellist* is "string" "string" [...]. The number of labels must equal the number of transitions.

`[x|y]line(`*numlist*`)` see: [G] **added line options**

`title(`*title*`)` see: [G] **title options**

`name(`*name* `[,` *replace*`])` see: [G] **name option**

`[y|x]scale(`*axis sub options*`)` see: [G] **axis scale options**

`[y|x]label(`*rule or values*`)` see: [G] **axis options**

`[y|x]size(`*#*`)` see: [G] **region options**

Syntax for `predict`

`predict` [`type`] *newvar* [*if*] [*in*] [`,` *statistic* `outcome(`*#*`)` `transition(`*#*`)` `choice(`*#*`)` `equation(`*#*`)`]

Options for `predict`

`transition(`*#*`)` specifies the transition, 1 is the first transition specified in the `tree` option in `seqlogit`, 2 the second, etc.

`choice(`*#*`)` specifies the choice within the transition, 0 is the first choice (the reference category), 1 the second, etc.

`equation(#)` specifies the equation, #1 is the first equation, #2 the second, etc. The "#" before the number is required.

statistic	description
`xb`	linear predictor
`stdp`	standard error of the linear predictor
`trpr`	probability of passing transition
`tratrisk`	proportion of respondents at risk of passing transition
`trvar`	variance of the indicator variable indicating whether or not the respondent passed the transition
`trgain`	difference in expected highest achieved level between those that pass the transition and those that do not
`trweight`	weight assigned to transition
`pr`	probability that an outcome is the highest achieved outcome.
`y`	expected highest achieved level
`effect`	Effect of variable of interest on expected highest achieved level. This variable is specified in the `ofinterest()` option in `seqlogit`. Interactions with the variables specified in the `over()` option of `seqlogit` are automatically taken into account.
`residuals`	difference between highest achieved level and expected highest achieved level.
`score`	first derivative of the log likelihood with respect to the linear predictor.

Syntax for `uhdesc`

`uhdesc` [, `at(`*atlist*`)` `overat(`*overatlist*`)` `ovarlab(`*stringlist*`)` `draws(#)`]

Options for `uhdesc`

`overat(`*overlist*`)` Specifies the values of the explanatory variables of the groups that are to be compared. It overrides any value specified in the `at()` option. Each comparison is separated by a comma. The syntax for *overlist* is:

varname_1 #[*varname_2* #[...]], *varname_1* #[*varname_2* #[...]], [...]

`at(`*atlist*`)` specifies the values at which the equations are evaluated. The syntax for *atlist* is: *varname*_1 #[*varname*_2 #...]. The equations will be evaluated at the mean values of any of the variables not specified in `at()` or `overat()`.

Say the dependent variable is highest achieved level of education, which is influenced by child's Socioeconomic Status (*ses*) and cohort (*coh*) and the interaction between *ses* and *coh* (*_ses_X_coh*). We want to compare the decomposition of the effect of *ses* over different cohorts for mean value of *ses*. Say that *coh* has only three values: 1, 2, and 3 and the mean value of *ses* is .5. Then the `overat()` and `at()` options would read:

```
overat( coh 1, coh 2, coh 3 ) at( ses .5 )
```

Notice that the values for the interaction term need not be specified in the `overat()` option, as long as it was created using the `over()` option in `seqlogit`.

`overlab(`*stringlist*`)` specifies the label that is to be attached to each group specified in the `overatlist()` option. Spaces are not allowed but an "_" will be displayed as an space. The number of labels has to be the same as the number of groups specified in the `overatlist()` option.

To continue the example above: say that a value of 1 on the variable `coh` corresponds to the cohort born in 1950, a value 2 corresponds to the cohort born in 1970, a value 3 corresponds to the cohort born in 1990, then the `overlab()` option would read:

```
overlab(1950 1970 1990)
```

`draws(`*#*`)` specifies the number of pseudo random draws from the distribution of the unobserved variable used for computing the descriptive statistics. These pseudo random draws are created using a Halton sequence, see: [M-5] **halton()**. `uhdesc` uses by default the same number of draws as specified in `seqlogit`.

Ongelijkheid in onderwijsuitkomsten en onderwijskansen in Nederland in de 20ste eeuw

Nederlandse samenvatting / Summary in Dutch

Kinderen van ouders met betere beroepen en een hogere opleiding krijgen over het algemeen een hogere opleiding dan kinderen van ouders met slechtere beroepen of een lagere opleiding. Deze samenhang tussen sociale herkomst en opleiding is zowel in Nederland (bijvoorbeeld, De Graaf en Ganzeboom 1993 of Gesthuizen en anderen 2005) als internationaal (bijvoorbeeld Shavit en Blossfeld 1993 of Breen en anderen 2009) gevonden. Gedeeltelijk hangt deze onderwijs ongelijkheid samen met het feit dat ouders met meer sociale, culturele, en economische middelen beter in staat zijn hun kinderen te helpen. Dit schept een paradox. Aan de ene kant is het goed dat ouders zoveel om hun kinderen geven dat ze alles doen om hun kinderen een voordeel te geven. Aan de andere kant willen we niet dat sommige kinderen een voordeel krijgen en andere een nadeel, alleen maar omdat sommigen de 'juiste' ouders hebben en andere niet. Een mogelijke oplossing van dit probleem is het onderwijssysteem zo in te richten dat de extra inspanningen van ouders minder invloed hebben. Het onderwijssysteem in een land kan op deze wijze een grote invloed uitoefenen op hoe personen op bevoorrechte posities belanden: vooral op basis van eigen vaardigheden en inzet of meer op basis van afkomst (Blau en Duncan 1967). Om deze reden kent het onderzoek naar de samenhang tussen sociale afkomst van kinderen en hun opleiding een lange geschiedenis. Recente overzichten van deze onderzoekstraditie zijn geschreven door Hout en DiPrete (2006) en Breen en Jonsson (2007).

Het doel van dit proefschrift is om aan dit onderzoek bij te dragen door aan te tonen hoe een aantal methodologische vernieuwingen het mogelijk maakt om met reeds bestaande data tot nieuwe inzichten te komen. Deze nieuwe inzichten hebben betrekking op de volgende vraag:

> In welke mate, hoe, en wanneer heeft de trend in Nederland naar minder ongelijkheid in onderwijskansen en onderwijsuitkomsten tussen personen die uit verschillende sociale milieus komen plaatsgevonden?

In deze vraag wordt onderscheid gemaakt tussen twee vormen van onderwijsongelijkheid:

1. de ongelijkheid in onderwijs-uitkomsten, waarmee ik de sterkte van de samenhang tussen sociaal milieu van de ouders en de hoogst behaalde opleiding van

hun kinderen bedoel. In dit proefschrift heb ik dit Inequality of Educational Outcome (IEOut) genoemd.

2. De ongelijkheid in onderwijs-kansen, waarmee de samenhang tussen het ouderlijk milieu en de kansen om van het ene onderwijs niveau naar het andere onderwijs niveau te gaan bedoel. In dit proefschrift heb ik dit Inequality of Educational Opportunity (IEOpp) genoemd.

IEOut is relevant wanneer men geïnteresseerd is in de mate waarin het onderwijssysteem als geheel gekenmerkt wordt door ongelijkheid, bijvoorbeeld omdat men wil weten hoe deze ongelijkheid in het onderwijssysteem doorwerkt in andere type ongelijkheden zoals succes op de arbeidsmarkt, het vinden van een partner, en gezondheid. IEOpp is relevant wanneer men wil weten welke fase in de onderwijscarrière gekenmerkt wordt door de grootste ongelijkheid, bijvoorbeeld omdat men wil weten hoe IEOut tot stand gekomen is, of waar in het onderwijssysteem ingegrepen moet worden om ongelijkheid te verminderen.

Het proefschrift bestaat uit een zestal hoofdstukken. Hieronder geef ik een korte beschrijving van elk hoofdstuk. Daarna ga ik op een aantal onderwerpen dieper in.

Korte beschrijving hoofdstukken

Als uitgangspunt dient hoofdstuk 2[1]. Dit is een replicatie van een studie door De Graaf en Ganzeboom (1993), die aangeeft wat men van de meest recente gegevens met de bestaande ‘standaard’ methoden leren kan. De overige hoofdstukken welke extra inzichten de nieuwe methoden opleveren. Hoofdstukken 3, 4, 5 laten drie manieren zien om de schatting van IEOut te verbeteren.

Hoofdstuk 3 neemt de meting van onderwijsniveau onder de loep. Traditioneel wordt bij de schatting van IEOut aan iedere opleiding een waarde toegekend op basis van ‘standaard jaren opleiding’. Met behulp van nieuw methoden wordt de waarde die aan iedere opleiding toegekend empirisch geschat op basis van de beroepsstatus die de respondenten met een gegeven opleiding verkregen hebben. Hieruit blijkt dat met name de waarde van lager beroeps onderwijs door de traditionele methode behoorlijk overschat wordt.

In hoofdstuk 4 is onderzocht hoe IEOut in de loop der jaren veranderde. Traditionele methoden hebben tot nog toe gevonden dat IEOut gestaag afneemt. Nieuw methoden stellen dit beeld bij: deze trend was met name een fenomeen in de jaren ’40 en ’50 voor mannen en in de jaren ’50 en ’60 voor vrouwen.

[1] Hoofdstuk 1 is de inleiding.

In hoofdstuk 5 is gekeken of het effect van de moeder op de opleiding van het kind over de tijd relatief belangrijker geworden is ten opzichte van het effect van de vader. Daarnaast is onderzocht of het effect van de opleiding van de ouders relatief belangrijker is geworden ten opzichte van het effect van de beroepsstatus. Uit dit hoofdstuk blijkt dat deze verhoudingen gedurende de onderzochte periode (1939–1991) onveranderd zijn gebleven.

In hoofdstuk 6 onderzocht ik de relatie tussen IEOpp en IEOut. Deze twee verschillende vormen van onderwijsongelijkheid zijn nauw met elkaar verbonden. IEOpp beschrijft ongelijkheden in het proces dat leidt tot een bepaald opleidingsniveau, terwijl IEOut de ongelijkheid in de uitkomst van dat proces beschrijft. Toch zijn beide vormen tot nog toe apart onderzocht.

Dit hoofdstuk beschrijft een nieuwe methode om hun onderlinge samenhang te beschrijven door gewichten te berekenen voor iedere overgang tussen onderwijsniveaus. Deze gewichten geven aan hoe belangrijk ongelijkheid gedurende iedere overgang (de IEOpps) is voor de ongelijkheid in het uiteindelijk behaalde onderwijsniveau (de IEOut). Hieruit blijkt dat aan het begin van de 20$^{\text{ste}}$ eeuw IEOut voornamelijk veroorzaakt werd door ongelijkheid gedurende de eerste transitie (of men na het basisonderwijs nog een diploma behaald of niet), terwijl voor recente cohorten de tweede transitie dominant is (of men een beroepsgerichte of academische richting op gaat).

In hoofdstuk 7 wordt nader ingegaan op een invloedrijke kritiek van Cameron en Heckman (1998) op het model dat het meest gebruikt wordt voor het schatten van IEOpp, de sequentiële logistische regressie. Hun kritiek is gebaseerd op de waarneming dat dit model extra gevoelig kan zijn voor vertekende invloeden van niet geobserveerde variabelen. De mogelijke sterkte van deze verstorende invloeden is in dit hoofdstuk onderzocht door modellen te schatten onder verschillende aannames over deze niet geobserveerde variabelen, en vervolgens te kijken hoe extreem deze aannames moeten zijn voordat de conclusies veranderen. Het resultaat was dat de kwalitatieve conclusies voor Nederland slechts onder zeer extreme aannames veranderden, maar dat de omvang van de IEOpps end de trends daarin waarschijnlijk onderschat worden. Dit betekend dat de problemen die door Cameron en Heckman werden aangekaart niet groot genoeg zijn de resultaten van statistische tests wezenlijk te beïnvloeden. Deze problemen hebben echter wel invloed op schatting de IEOpps en de trends daarin. Ondanks dat blijven deze schatting nog steeds bruikbaar wanneer ze geïnterpreteerd worden als een “ondergrens”, dat wil zeggen, de werkelijke waarden voor de IEOpps en hun trends liggen naar alle waarschijnlijkheid niet onder deze schatting.

Replicatie

De dissertatie begint met een replicatie van de Nederlandse bijdrage van De Graaf en Ganzeboom (1993) aan een invloedrijke internationale vergelijking door Shavit en Blossfeld (1993). Deze replicatie creeert een referentiepunt dat weergeeft wat met de 'standaard' methoden uit recente data geleerd kan worden. De Graaf en Ganzeboom (1993) bestudeerden zowel IEOpp als IEOut wat gezien de belangrijke rol van beide begrippen in dit proefschrift een voordeel is. Bovendien, zijn de data die in dit proefschrift gebruikt wordt een uitbreiding van de data die De Graaf en Ganzeboom (1993) gebruikten. De originele data bestonden uit 10 enquêtes die zijn samengevoegd tot een gegevensbestand. Later hebben Ganzeboom en Treiman (2009) dit gegevensbestand in het kader van het 'International Stratification and Mobility File' (ISMF) uitgebreid door meer enquêtes toe te voegen. Het Nederlandse deel van de ISMF bestond ten tijde van dit proefschrift uit 54 enquêtes. De replicatie gebruikte hierdoor ongeveer 5 maal meer respondenten dan de originele studie (69.868 respondent versus 11.244 respondenten) en omvatte 20 extra jaren (1891–1980 versus 1891–1960).

De belangrijkste bevinding van deze replicatie is dat, ondanks deze veel uitgebreidere data, de resultaten in grote lijnen overeenkomen met de resultaten van de oorspronkelijke studie. Met de standaard methoden zijn de volgende trends in IEOpp en IEOut gevonden. Voor IEOut geldt een significante negatieve trend. De trend in IEOpp verschilt per transitie. Voor de overgang van lager onderwijs naar het behalen van een vervolg diploma werd een significante neergaande trend gevonden. Voor de doorstroom van LBO en MAVO[2] naar hogere opleidingsniveaus werd in een aantal gevallen een significant negatieve trend gevonden terwijl in andere gevallen geen significante trend werd gevonden. Voor de transition van HAVO, VWO, en MBO aan de ene kant naar HBO en WO aan de andere kant werd in de meeste gevallen geen trend gevonden en in een aantal gevallen een positieve trend. De meeste van de trends in IEOpp en IEOut waren linear.

IEOut: operationalisatie van opleiding, trends en hulpbronnen van families

In hoofdstukken 3, 4, en 5 ligt de nadruk op het schatten van de trend in IEOut.

Een van de zwakke punten van de wijze waarop IEOut geschat werd in hoofdstuk 2 is dat, in navolging van De Graaf en Ganzeboom (1993), opleidingsniveau werd uitgedrukt in 4 opleidingscategorieën die de waardes 1 tot 4 kregen. Hierdoor wordt

[2] Ik gebruik de namen van na de Mammoet Wet, maar bedoel daarmee ook de equivalente niveaus van voor de Mammoet Wet.

impliciet verondersteld dat de afstanden tussen de verschillende categorieën gelijk zijn. Een populair alternatief is de opleidingen een waarde te geven op basis van het aantal jaren dat een 'standaard' student nodig heeft om dat niveau te bereiken. Een nadeel van deze methode is dat er vaak *ad hoc* aanpassingen nodig zijn om te zorgen dat de rangorde overeenkomt met wat *a priori* bekend is over de opleidingen. Een veel gebruikte schaal van dit type is de *a priori* schaal van Ganzeboom en Treiman (2009).

Hoofdstuk 3 verbeterde deze standaard manieren van toekennen van waarden aan de opleidingscategorieën door deze waarden empirisch te schatten, zodanig dat de resulterende opleidingsschaal optimaal is voor het voorspellen van beroepsstatus. Deze geschatte schaal werd vervolgens vergeleken met de *a priori* schaal van Ganzeboom en Treiman (2009). De empirische schaal komt grotendeels overeen met de *a priori* schaal, met als belangrijkste uitzondering dat de waarde van het LBO in de *a priori* schaal overschat werd. Dit betekent dat respondenten met LBO gemiddeld een aanzienlijk lagere beroepsstatus hadden dan was voorspeld met behulp van de *a priori* schaal. De resultaten gaven bovendien aan dat de veranderingen in de arbeidsmarkt gedurende de onderzochte periode (1958 tot 2006) weinig effect gehad hebben op de relatieve afstanden tussen onderwijs categorieën. Daarentegen hebben, ten tijde van de invoering van "Mammoet Wet", de MAVO en het HBO relatief aan waarde verloren. De verandering in de waarde van de MAVO was te verwachten, aangezien dit niveau veranderde van een niveau dat voorbereid voor de arbeidsmarkt tot een niveau dat voorbereidt op een volgend niveau van het onderwijs (MBO). Een mogelijke reden voor de verandering in de waarde in het HBO zou kunnen zijn dat het toegankelijk werd vanuit het toenmalig nieuwe niveau HAVO.

In hoofdstuk 4 werd onderzocht of de trend in IEOut is veranderd over de tijd. In eerder onderzoek is voornamelijk een constante negatieve trend in IEOut gevonden. Het is echter onwaarschijnlijk dat deze lineaire trend zich zal voortzetten, omdat dat uiteindelijk zou leiden to een negatieve samenhang tussen familie achtergrond en opleidingsniveau. De negatieve trend in IEOut zal dus op enig moment moeten vertragen en het doel van Hoofdstuk 4 was te proberen deze vertraging van de trend te ontdekken. Ik heb inderdaad bewijzen gevonden dat de trend niet-lineair is, maar de verwachte vertraging in de dalende trend in de IEOut is niet gevonden. Voor zowel mannen als vrouwen is een periode van negatieve trend gevonden (respectievelijk 1941 – 1960 en 1952 – 1977). Er is dus geen significante trend gevonden voor recente cohort (mannen die 12 jaar oud waren na 1960 en vrouwen die 12 jaar oud waren na 1977). Bij dergelijke schattingen is het statistisch onderscheidingsvermogen het geringst bij de jongste en oudste cohorten, dus de afwezigheid van een significante trend kan ook daardoor verklaard worden. Alleen voor mannen zijn er enige aanwijzingen gevonden dat de periode van niet-significante trend is voorafgegaan door een vertraging, maar deze vertraging is (nog) niet significant. Er zijn wel duidelijke aanwijzingen dat de dalende trend werd voorafgegaan door een periode waarin de trend aanzienlijk versnelde (1935 – 1944 voor mannen en 1949 – 1952 voor vrouwen).

De reden dat veranderingen in de trend gevonden werden terwijl eerder onderzoek deze niet kon waarnemen, is een verschil in de methoden die worden gebruikt bij het testen voor niet-lineariteiten. De standaard methoden bestonden uit een schatting van een niet-lineaire trend met behulp van een kwadratische of discrete trend. De kwadratische trend is vaak niet flexibel genoeg om eventuele niet-lineariteiten in de trend te kunnen waarnemen. De discrete trend is vaak juist te flexibel, waardoor teveel statistische onderscheidingsvermogen verloren gaat. Als alternatief is de trend geschat met behulp van een lokale polynomiale curve. Deze is flexibeler dan een kwadratische curve maar behoudt meer statistisch onderscheidingsvermogen dan een discrete curve.

Ik heb ook gekeken of de in hoofdstuk 3 geschatte schaal tot andere conclusies leidt dan de *a priori* schaal van Ganzeboom en Treiman (2009). De geschatte schaal voor opleiding leidde tot een iets stabielere trend (minder extreme uitschieters bij de jongste en oudste cohorten) dan de *a priori* schaal van Ganzeboom en Treiman (2009).

In hoofdstuk 5 heb ik gekeken welke ouder de meeste invloed had op de opleiding van de kinderen. Hierbij heb ik niet alleen onderscheid gemaakt tussen de vader en de moeder, maar ook tussen de ouder met de hoogste status en de ouder met de laagste status, en tussen de ouder met hetzelfde geslacht als het kind en de ouder van het andere geslacht. Bovendien heb ik gekeken naar de relatieve invloed van de beroepsstatus en de hoogst behaalde opleiding van de ouders.

Met betrekking tot welk type onderscheid tussen ouders de meeste invloed uitoefend heb ik gevonden dat het onderscheid tussen de ouder met de hoogste status en de ouder met de laagste status belangrijker is dan het onderscheid tussen de vaders en moeders of het onderscheid tussen de ouder met hetzelfde geslacht als het kind en de ouder van het andere geslacht. Met betrekking tot welke inbreng van de ouders het belangrijkste is heb ik matig bewijs gevonden dat beroepsstatus belangrijker is dan de opleiding van de ouders. Daarnaast vond ik dat het thuisblijven van de moeder alleen een negatief effect heeft op het opleidingsniveau van de kinderen als de moeder weinig onderwijs heeft *en* de vader heeft een baan heeft met een lage status. Dit effect wordt echter positief als de moeder goed opgeleid is of wanneer de vader een baan met een hoge status heeft.

Daarnaast heb ik ook onderzocht of deze patronen veranderd zijn over de tijd. Ik had verwacht dat veranderingen in de rollen van mannen en vrouwen in de samenleving gedurende de onderzochte periode (1939 tot 1991) ook zou leiden tot een verandering in de verhouding de effecten van de moeder en de effecten van de vader. Ook verwachtte ik dat het effect van beroepsstatus zou dalen ten opzichte van het effect van opleiding. De achterliggende redenering is dat beroepsstatus meer verbonden is met de economische middelen die beschikbaar zijn in het gezin, en dat economische beperkingen over de tijd minder invloed op opleiding zouden krijgen. Deze afname van de invloed van economische beperkingen in een gezin komt aan de ene kant doordat de grote economische groei ervoor gezorgt heeft dat bijna iedereen welvarender geworden is en aan de andere kant doordat het onderwijs zwaarder wordt gesubsidieerd. Om deze hypothesen te toetsen, heb ik gebruik gemaakt van een model met parametrisch gewogen covariaten zoals voorgesteld door Yamaguchi (2002). Dit model schat de effecten onder de nulhypothese dat de relatieve effecten van de opleiding en beroepsstatus van beide ouders onveranderd zijn gebleven over de tijd. In tegenstelling tot wat ik verwachtte, kan deze hypothese niet worden afgewezen.

De relatie tussen IEOpp en IEOut

Bij het onderzoek naar ongelijkheid in toegang tot onderwijs is het nuttig om onderscheid te maken tussen ongelijkheid gedurende het onderwijsproces (de IEOpp) en de ongelijkheid in de uiteindelijke uitkomst van dat proces (de IEOut). Daarnaast is het ook goed verdedigbaar dat IEOpp en IEOut elkaar aanvullende informatie bevatten; een beschrijving van het onderwijsproces kan zinvol worden aangevuld met een beschrijving van de uitkomst van dat proces. Hoofdstuk 6 bespreekt een nieuwe methode die een geïntegreerde discussie van IEOpp en IEOut mogelijk maakt. Deze methode begint met het standaard model voor het schatten van IEOpps, het sequentiële logit model zoals voorgesteld door Mare (1981). De IEOpps die in dit model geschat worden zijn het effect van familie achtergrond op de waarschijnlijkheid dat iemand de overgang naar een volgend (hoger) onderwijsniveau maakt. In dit hoofdstuk toon ik aan dat dit model een decompositie van IEOut impliceert als een gewogen som van de IEOpps. Met andere woorden, de ongelijkheid in onderwijsuitkomsten is de som de ongelijkheden gedurende iedere stap in het onderwijsproces, maar niet iedere stap is even belangrijk. De “belangrijkheid” van iedere stap wordt weergegeven door een gewicht dat aan die stap wordt toegekend. Deze gewichten blijken het product te zijn van drie elementen:

1. het percentage van respondenten dat de overgang kan maken, waardoor een overgang meer gewicht krijgt wanneer hij meer mensen treft,

2. het percentage respondent dat slaagt voor de overgang maal het percentage respondenten dat niet slaagt voor deze overgang. Hierdoor krijgt een overgang minder gewicht wanneer ofwel vrijwel iedereen slaagt ofwel vrijwel iedereen niet slaagt, en

3. de verwachte toename in het hoogst bereikte onderwijsniveau als gevolg van het slagen bij een bepaalde stap, waardoopr een overgang meer gewicht krijgt naarmate respondent die slagen daar meer profijt van hebben.

Deze drie elementen maken het mogelijk schattingen van IEOpps aan te vullen met schattingen van hoe relevant deze zijn voor IEOut. Bovendien biedt deze decompositie een inhoudelijk interpreteerbaar mechanisme waardoor de toename in het gemiddelde onderwijsniveau onderwijsongelijkheid kan beïnvloeden. De toename van het gemiddelde onderwijsniveau hangt samen met een toename in de waarschijnlijkheden om te slagen bij de verschillende overgangen, waardoor de gewichten veranderen, wat op zijn beurt weer leidt tot veranderingen in IEOut. Dit is van belang omdat veel van de methoden die in eerder onderzoek gebruikt werden direct voor de toename in gemiddeld onderwijsniveau controleerden, waardoor het effect van de ontwikkeling niet onderzocht kon worden.

Bij de toepassing van deze decompositie op Nederland heb ik het Nederlands onderwijs systeem samengevat door onderscheid te maken tussen vier overgangen: De eerste overgang maakt onderscheid tussen diegene die vertrekken uit het onderwijs met alleen een diploma primair onderwijs en diegene die een hoger diploma behalen. De tweede overgang is toegankelijk voor diegene die doorgaan in het onderwijs, en maakt onderscheid tussen een 'beroepsgericht' pad (LBO en MAVO) en een 'academisch' pad (HAVO en VWO). De derde overgang is alleen toegankelijk voor diegene die het beroepsgerichte pad hebben gekozen en maakt onderscheid tussen diegene die vertrekken met alleen een LBO of MAVO diploma en diegene die een MBO diploma behalen. De vierde overgang is alleen toegankelijk voor diegene die het academische pad hebben gekozen en maakt onderscheid tussen diegene die alleen een HAVO of VWO diploma halen en diegene die ook nog een HBO of universitair diploma behalen.

Ik vond dat het merendeel van de IEOut veroorzaakt wordt door de eerste twee overgangen, en dat de laatste twee overgangen slechts een zeer klein deel van de IEOut verklaarden. Bovendien vond ik dat IEOut in het begin van de onderzochte periode (ongeveer 1905–1940) voornamelijk werd bepaald door de eerst transitie, terwijl de tweede transitie dominant is in recentere cohorten (ongeveer 1960–1990). De eerste overgang daalde snel in belang doordat het passeren van deze overgang bijna universeel werd. Diegene die niet voor deze overgang slagen komen nog steeds disproportioneel uit minder bevoorrechte milieus, maar het aantal personen dat niet slaagt is in recentere cohorten zo laag dat dit nauwelijks nog invloed heeft op IEOut. De tweede overgang is daarentegen sterk in belang toegenomen doordat meer mensen toegang hebben gekregen tot deze overgang en doordat van deze mensen nu een groter aandeel in het academische pad terechtkomt. Veranderingen in de ongelijkheid in onderwijsuitkomsten zijn dus niet zozeer opgetreden doordat het onderwijsproces 'eerlijker' geworden is. De verklaring ligt voornamelijk in het feit dat de eerste transitie, die gekenmerkt wordt door een zeer hoge sociale ongelijkheid, vervangen is door de minder ongelijke tweede transitie als dominante bron van IEOut.

IEOpp: de invloed van niet-waargenomen variabelen

In hoofdstuk 7 ga ik in op een invloedrijke kritiek van Cameron en Heckman (1998) op het standaard model voor het schatten van IEOpps, het sequentiële logit model. Cameron en Heckman (1998) betogen dat het sequentiële logit model, net als ieder ander model, niet alle variabelen bevat die de afhankelijke variabele beïnvloeden. Echter, deze niet geobserveerde variabelen kunnen een grotere invloed hebben dan gebruikelijk, voornamelijk omdat zij de resultaten kunnen beïnvloeden zelfs wanneer deze niet-geobserveerde variabelen oorspronkelijk ongecorreleerd zijn met de geobserveerde variabelen.

Dit suggereert dat men voor deze niet geobserveerde variabelen zou moeten controleren, maar dat is per definitie onmogelijk. Het is echter wel mogelijk een scenario te creëren over de niet-waargenomen variabelen en vervolgens de effecten te schatten gegeven dat scenario. In hoofdstuk 7 wordt een set van scenario's voorgesteld die nuttig kunnen zijn om de gevoeligheid van de schattingen voor niet-geobserveerde heterogeniteit te beoordelen. Bovendien wordt een methode voor het schatten van de effecten binnen deze scenario's besproken. Deze aanpak wordt geïllustreerd door middel van een replicatie van de analyse uit hoofdstuk 2. Hierbij werd gekeken naar de robuustheid van twee testen — of de IEOpps veranderen over cohorten en over transities. Daarnaast werd de robuustheid van de schattingen van de omvang van de IEOpps en de trend in IEOpps onderzocht. Uit de gevoeligheidsanalyse blijkt dat de resultaten van de statistische tests slechts veranderde in zeer extreme scenario's. De IEOpps en de trend in IEOpps namen echter al toe in gematigde scenario's, wat aangeeft dat modellen die niet voor niet-geobserveerde heterogeniteit controleren deze effecten waarschijnlijk onderschatten. In gematigde scenario's dalen de IEOps minder over transities dan in modellen die niet voor niet-geobserveerde heterogeniteit controleren. Dit betekent dat het algemeen gevonden patroon van afnemende effecten van de familie achtergrond variabelen over transities ten minste gedeeltelijk is te wijten aan niet-geobserveerde heterogeniteit.

Conclusies

De onderzoeksvraag van dit proefschrift is: "In welke mate, hoe, en wanneer heeft de trend in Nederland naar minder ongelijkheid in onderwijskansen en onderwijsuitkomsten tussen personen die uit verschillende sociale milieus komen plaatsgevonden?". Het antwoord is opgedeeld in de volgende elementen:

IEOut

- Er was een neerwaartse trend in IEOut in Nederland gedurende de jaren ’40 en ’50 voor mannen en de jaren ’50 en ’60 voor vrouwen. Dit heeft geleid to ongeveer een halvering van IEOut. Hieraan ging een periode van versnelling vooraf, en er zijn zelfs enige aanwijzingen dat de trend aanvankelijk stijgend was. Deze uitkomst is nieuw, aangezien eerdere studies de hypothese van een lineaire trend niet konden verwerpen.
- In deze dissertatie is een betere schaal voor de opleidingscategorieën geschat, maar deze nieuwe schaal had slechts een beperkt effect op de geschatte trend in IEOut.
- De relatieve invloed van de vader op de opleiding van zijn kinderen ten opzichte van het effect van de moeder op de opleiding van de kinderen bleef onveranderd. Dit geldt ook voor de relatieve invloed van de beroepsstatus van de ouders op de opleiding van hun kinderen ten opzichte van de invloed van de opleiding van de ouders op de opleiding van hun kinderen.

IEOut and IEOpp

- De ongelijkheid in IEOout trad aanvankelijk vooral op tijdens de transitie na het behalen van een diploma primair onderwijs. Het ging dan om uitstromen of verder leren. Deze transitie heeft veel aan belang ingeboet doordat tegenwoordig het overgrote merendeel in deze transitie slaagt. Hierdoor is ook de daaropvolgende transitie, die onderscheid maakt tussen een ‘beroepsgericht’ pad (LBO en MAVO) en een ‘academisch’ pad (HAVO en VWO), belangrijker geworden.
- Deze verschuiving verklaart zowel de aanvankelijke stijging in IEOut als de latere daling van IEOut. De oorzaak van de stijging is het feit dat de daling in het belang van de eerste transitie ruim gecompenseerd werd door een toename in het belang van de tweede transitie. De daling was het gevolg van het feit dat de minder ongelijke tweede transitie de eerste transitie vrijwel kompleet vervangen heeft als de dominante bron van IEOut.

IEOpp

- Er zijn significant neergaande trends in IEOpps gevonden voor de eerste transities van de onderwijscarrière. Voor latere transities zijn significant negatieve, niet significante, en significant positieve trends in IEOpps gevonden.
- De IEOpps voor de eerste transities van de onderwijscarrière zijn groter dan de IEOpps voor latere transities.
- Een gevoeligheidsanalyse heeft aangetoond dat deze conclusies in kwalitatieve zin robuust zijn, maar dat de omvang van de IEOpps en de trends daarin waarschijnlijk onderschat worden door modellen die geen rekening houden met niet geobserveerde heterogeniteit.

Data references

1, net58 Gadourek, I. (1958) ‘Health Threatening Habits’, The Hague: DANS, nr. P0142 [distributor].

2, net67 Kooij, G.A, and J.W. te Kloeze (1967) Family in Modern City Environment 1967, The Hague: DANS, nr. P0463 [distributor].

3, net67t Galtung, J. (1967) ‘Images of the year 2000’, The Hague: DANS, nr. P0042 [distributor].

4, net70 Stouthard, P.C., F.J. Heunks, J.W.W. Thomassen, W.E. Miller, and J. Rusk (1970) National Election Study, The Hague: DANS, nr. P0136 [distributor].

5, net71c Inglehart, R., and Rabier, J-R (1971) ‘European Communities Study’, Ann Arbor, MI.: Inter-university Consortium for Political and Social Research, nr. 7275 [distributor]

6, net71 Verba, S., Nie, N.H., and Kim, J. (1971) ‘Political Participation and Equality in Seven Nations’, Ann Arbor, MI: Inter-University Consortium for Political and Social Research, nr 7768 [distributor].

7, net74p Barnes, S.H., Kaase, M., et al. (1974) ‘Political Action: An Eight Nation Study’, Ann Arbor, MI: Inter-University Consortium for Political and Social Research, nr. 7777 [distributor].

8, net76j Hermkens, P.L.J., and Wijngaarden, P.J. van (1976) ‘Criteria for Justification of Income Differences Netherlands’, The Hague: DANS, nr. P0653 [distributor].

9, net77 CBS (Centraal Bureau voor de Statistiek) (1977) ‘Life [Situation] Survey, Netherlands’, The Hague: DANS, nr. P0328 [distributor].

10, net77e Irwin, G.A., J. Verhoef, and C.J. Wiebrens (1977) ‘Dutch Parliamentary Election Study’, The Hague: DANS, nr. P0354 [distributor].

11, net79p Heunks, F.J., P.C. Stouthard, J.W. van Deth, and J.J.A. Thomassen (1979) ‘Political Action II, Nederland’, The Hague: DANS, nr. P0823 [distributor].

12, net81e Eijk, C. van der, B. Niemoeller, and A.Th.J. Eggen (1981) ‘Dutch Parliamentary Election Study’, The Hague: DANS, nr. P0350 [distributor].

13, net82e Eijk, C. van der, M.J. Koopman, B. Niemoeller, and S.M.J. Koopman (1982) 'Dutch Parliamentary Election Study', The Hague: DANS, nr. P0633 [distributor].

14, net82n Heinen, A., and A. Maas (1982) 'NPAO Labour Market Survey, Netherlands', The Hague: DANS, nr. P0748 [distributor].

15, net82u Ultee, W.C., and H. Sixma (1982) 'National Prestige Survey, Netherlands', The Hague: DANS, nr. P0839 [distributor].

16, net85o Organisatie voor Strategisch Arbeidsmarktonderzoek (OSA) (1985) 'National Labour Market Survey', The Hague: DANS, nr. P1462 [distributor].

17, net86e Eijk, C. van der, B. Niemoeller, and G.A. Irwin (1986) 'Dutch Parliamentary Election Study', The Hague: DANS, nr. P0979 [distributor].

18, net86l CBS (Centraal Bureau voor de Statistiek) (1986) 'Life [Situation] Survey', The Hague: DANS, nr. P1454 [distributor].

19, net87i Becker, J.W., and A.N. Nauta (1987) 'International Social Survey 1987', Ann Arbor, MI: Inter-University Consortium for Political and Social Research, nr. 9025 [distributor].

20, net87j Hermkens, P.L.J., and P.J. van Wijngaarden (1987) 'Criteria for justification of income differences 1976', The Hague: DANS, nr. P1194 [distributor].

21, net87s Fiselier, A.A.M., M.G.M. van der Poel, and A. Felling (1987) 'Primary relations and social support 1987', The Hague: DANS, nr. P1172 [distributor].

22, net88o Organisatie voor Strategisch Arbeidsmarktonderzoek (OSA) (1988) 'National Labour Market Survey', The Hague: DANS, nr. P1462 [distributor].

23, net90 Eisinga, R., A. Felling, J. Peters, P.L.H. Scheepers, and O. Schreuder (1990) 'Religion in Dutch Society, 1990 – SOCON'90', The Hague: DANS, nr. P1100 [distributor].

24, net90o Organisatie voor Strategisch Arbeidsmarktonderzoek (OSA) (1990) 'National Labour Market Survey', The Hague: DANS, nr. P1462 [distributor].

25, net91j Alwin, D.F., D.M. Kluegel, and M. Dielman (1991) 'International Social Justice Project', Ann Arbor, MI: Inter-University Consortium for Political and Social Research, nr. 6705 [distributor].

26, net92f Ultee, W.C., and H.B.G. Ganzeboom (1992) 'Netherlands Family Survey 1992–1993', The Hague: DANS, nr. P1245 [distributor].

27, net92o Organisatie voor Strategisch Arbeidsmarktonderzoek (OSA) (1992) 'National Labour Market Survey', The Hague: DANS, nr. P1462 [distributor].

28, net92t Stichting Sociaal-culturele Wetenschappen (1992) 'Aspects of life-event history of the Dutch population', The Hague: DANS, nr. P1107 [distributor].

29, net94e Anker, H. and E.V. Oppenhuis (1994) 'Dutch Parliamentary Election Study', The Hague: DANS, nr. P1208 [distributor].

30, net94h Weesie, J., and H.B.G. Ganzeboom (1994) 'Households in the Netherlands 1994', The Hague: DANS, nr. P1306 [distributor].

31, net94o Organisatie voor Strategisch Arbeidsmarktonderzoek (OSA) (1994) 'National Labour Market Survey', The Hague: DANS, nr. P1462 [distributor].

32, net95h Weesie, J., M. Kalmijn, and H.B.G. Ganzeboom (1995) 'Households in the Netherlands, 1995', The Hague: DANS, nr. P1458 [distributor].

33, net95s Eisinga, R., A. Felling, and R. König (1995) 'Religion in Dutch Society, 1995', The Hague: DANS, nr. P1336 [distributor].

34, net95y Weesie, J., M. Kalmijn, and H.B.G. Ganzeboom (1995) 'Households in the Netherlands, 1995' [producer], The Hague: DANS, nr. P1458 [distributor].

35, net96 Gijsberts, M. and Ganzeboom, H.B.G. (1996) 'Social Inequality in the Netherlands 1996', Groningen, Nijmegen, Utrecht: Interuniversity Center for Social Science Theory and Methodology.

36, net96c Wittebrood, K., and M.J. ter Voert, (1996) 'Netherlands Survey on Criminality and Law Enforcement 1996', The Hague: DANS, nr. P1465 [distributor].

37, net96o Organisatie voor Strategisch Arbeidsmarktonderzoek (OSA) (1996) 'National Labour Market Survey', The Hague: DANS, nr. P1462 [distributor].

38, net96y Wittebrood, K., and M.J. ter Voert, (1996) 'Netherlands Survey on Criminality and Law Enforcement 1996, Youth Sample', The Hague: DANS, nr. P1465 [distributor].

39, net98 Nieuwbeerta, P. (1998) 'Social and Economic Attitudes in the Netherlands 1998', The Hague: DANS, nr. P1435 [distributor].

40, net98e Aarts, K., Kolk, H. van der, and M. Kamp (1998) 'Dutch Parliamentary Election Study', The Hague: DANS, nr. P1415 [distributor].

41, net98f Graaf, N.D. de, P.M. de Graaf, G. Kraaykamp, and W.C. Ultee (1998) ‘Family Survey Dutch Population 1998’, The Hague: DANS, nr. P1583 [distributor].

42, net98o Organisatie voor Strategisch Arbeidsmarktonderzoek (OSA) (1998) ‘National Labour Market Survey’, The Hague: DANS, nr. P1462 [distributor].

43, net99 Dijk, L. van, H.B.G. Ganzeboom, J. de Haan, and S. Rijken (1999) ‘Use of new communication sources, 1999’, The Hague: DANS, nr. P1571 [distributor].

44, net99a Sociaal en Cultureel Planbureau (SCP) (1999) ‘Amenities and Services Utilization Survey, 1999’, The Hague: DANS, nr. P1513 [distributor].

45, net99i International Social Survey Programme (1999) ‘International Social Survey 1999’, Cologne: GESIS, nr. 3430 [distributor].

46, net00f Graaf, P.M. de, N.D. de Graaf, G.L.M. Kraaykamp, and W.C. Ultee (2000) ‘Family Survey Dutch Population 2000’, The Hague: DANS, nr. P1609 [distributor].

47, net00s Eisinga, R.N., M.T.A. Coenders, A.J.A. Felling, M. te Grotenhuis, S. Oomens, and P.L.H. Scheepers (2000) ‘Religion in Dutch Society, 2000’, The Hague: DANS, nr. P1556 [distributor].

48, net02e Irwin, G.A., J.J.M. Holsteyn, and J.M. den Ridder (2002) ‘Dutch Parliamentary Election Study’, The Hague: DANS, nr. P1628 [distributor].

49, net03f Graaf, N.D. de, P.M. de Graaf, G.L.M. Kraaykamp, and W.C. Ultee (2003) ‘Family Survey Dutch Population 2003’.

50, net03n Dykstra, P.A., M. Kalmijn, G.C.M. Knijn, A.E. Komter, C.H. Mulder, and A.C. Liefbroer (2003) ‘Netherlands Kinship Panel Study - wave 1’, The Hague: DANS, nr. P1635 [distributor].

51, net04e ‘European Social Survey’ (2004) ‘European Social Survey’, Bergen: Norwegian Social Science Data Services [distributor].

52, net04i International Social Survey Programme (2004) ‘International Social Survey 2004’, Cologne: GESIS, nr. 3950 [distributor].

53, net06e Aarts, K, H. van der Kolk, and J.J.A. Thomassen (2006) ‘Dutch Parliamentary Election Study’, The Hague: DANS, nr. P1719 [distributor]

54, net06i Ganzeboom, H.B.G., and Opdam, S. (2006) ‘International Social Survey 2006’.

Bibliography

Acock, Alan C. and Wen Shan Yang. 1984. "Parental power and adolescents' identification." *Journal of Marriage and the Family* 46:487–495.

Agresti, Alan. 2002. *Categorical Data Analysis*. Hoboken, NJ: Wiley-Interscience, 2nd edition.

Allison, Paul D. 1999. "Comparing Logit and Probit Coefficients Across Groups." *Sociological Methods & Research* 28:186–208.

Allison, Paul D. 2002. *Missing Data*. Quantitative Applications in the Social Sciences, 136. Thousand Oaks: Sage.

Arends-Kuenning, Mary and Suzanne Duryea. 2006. "The effect of parental presence, parents' education, and household headship on adolescents' schooling and work in Latin America." *Journal of Family and Economic Issues* 27:263–286.

Bakker, Bart F.M. and Peter G.J. Cremers. 1994. "Gelijke Kansen in het Onderwijs? Een Vergelijking van Vier Cohorten Leerlingen in hun Overgang naar het Voortgezet Onderwijs." *Tijdschrift voor Onderwijsresearch* 19:191–203.

Bakker, Bart F.M., Jaap Dronkers, and Huibert Schijf. 1982. "Verandering in Individuele Schoolloopbanen tussen 1959 en 1977 in de Stad Groningen." *Mens en Maatschappij* 57:253–266.

Bakker, Bart F.M. and Siebe P. Schouten. 1991. "Trends in Onderwijskansen: Een Vergelijking van de Overgang naar het Voortgezet Onderwijs van de Generaties Geboren rond 1953, 1965, en 1971." *Sociale Wetenschappen* 34:1–21.

Blau, Peter M. and Otis D. Duncan. 1967. *The American Occupational Structure*. New York: Wiley.

Bollen, Kenneth A. 1984. "Multiple Indicators: Internal Consistency or No Necessary Relationship?" *Quality & Quantity* 18:399–385.

Bollen, Kenneth A. and Richard Lennox. 1991. "Conventional Wisdom on Measurement: A Structural Equation Perspective." *Psychological Bulletin* 110:305–314.

Boudon, Raymond. 1974. *Education, Opportunity and Social Inequality*. New York: Wiley.

Boyd, Carol J. 1989. "Mothers and Daughters: A Discussion of Theory and Research." *Journal of Marriage and the Family* 51:291–301.

Breen, Richard. 1996. *Regression Models: Censored, Sample Selected, or Truncated Data*. Quantitative Applications in the Social Sciences, 111. Thousand Oaks: Sage.

Breen, Richard and Jan O. Jonsson. 2000. "Analyzing Educational Careers: A Multinomial Transition Model." *American Sociological Review* 65:754–772.

Breen, Richard and Jan O. Jonsson. 2005. "Inequality of Opportunity in Comparative Perspective: Recent Research on Educational and Social Mobility." *Annual Review of Sociology* 31:223–243.

Breen, Richard and Jan O. Jonsson. 2007. "Explaining Change in Social Fluidity: Educational Equalization and Educational Expansion in Twentieth-Century Sweden." *American Journal of Sociology* 112:1775–1810.

Breen, Richard, Ruud Luijkx, Walter Müller, and Reinhard Pollak. 2009. "Nonpersistent inequality in educational attainment: Evidence from eight European countries." *American Journal of Sociology* 114:1475–1521.

Buis, Maarten L. 2007a. "`propcnsreg`: Stata Module Fitting a Measurement Model with Causal Indicators." `http://econpapers.repec.org/software/bocbocode/s456858.htm`.

Buis, Maarten L. 2007b. "`seqlogit`: Stata Module to Fit a Sequential Logit Model." `http://econpapers.repec.org/software/bocbocode/S456843.htm`.

Buis, Maarten L. 2009a. "`proprcspline`: Stata module for restricted cubic spline smoothing of proportions." `http://econpapers.repec.org/software/bocbocode/s457062.htm`.

Buis, Maarten L. 2009b. "`sheafcoef`: Stata Module to Compute Sheaf Coefficients." `http://econpapers.repec.org/software/bocbocode/s456945.htm`.

Cameron, Stephen V. and James J. Heckman. 1998. "Life Cycle Schooling and Dynamic Selection Bias: Models and Evidence for Five Cohorts of American Males." *The Journal of Political Economy* 106:262–333.

Chevalier, Arnaud and Gauthier Lanot. 2002. "The relative effect of family characteristics and financial situation on educational achievement." *Education Economics* 10:165–181.

Cleveland, William S. 1979. "Robust Locally Weighted Regression and Smoothing Scatterplots." *Journal of the American Statistical Association* 74:829–836.

Clogg, Clifford C. and James W. Shockey. 1984. "Mismatch Between Occupation and Schooling: A Prevalence Measure, Recent Trends and Demographic Analysis." *Demography* 21:235–257.

Colding, Bjørg. 2006. "A dynamic analysis of educational progression of children of immigrants." *Labour Economics* 13:479–492.

Cragg, John G. and Russell S. Uhler. 1970. "The Demand for Automobiles." *The Canadian Journal of Economics* 3:386–406.

De Graaf, Nan Dirk, Paul M. De Graaf, and Gerbert Kraaykamp. 2000. "Parental Cultural Capital and Educational Attainment in the Netherlands: A refinement of the Cultural Capital Perspective." *Sociology of Education* 73:92–111.

De Graaf, Paul M. and Harry B.G. Ganzeboom. 1990. "Intergenerational Educational Mobility in the Netherlands for Birth Cohorts from 1891 through 1960." *Netherlands' Journal of Social Sciences* 26:35–50.

De Graaf, Paul M. and Harry B.G. Ganzeboom. 1993. "Family Background and Educational Attainment in the Netherlands for the 1891-1960 Birth Cohorts." In *Persistent Inequality: Changing Educational Attainment in Thirteen Countries*, edited by Yossi Shavit and Hans-Peter Blossfeld, chapter 4, pp. 75–99. Boulder: Westview Press.

De Graaf, Paul M. and Ruud Luijkx. 1992. "Van 'Ascriptie' naar 'Achievement'? Trends in Statusverwerving in Nederland tussen 1930 en 1980." *Mens en Maatschappij* 67:412–433.

De Graaf, Paul M. and Ruud Luijkx. 1995. "Paden naar Succes: Geboorte of Diploma's?" In *Verschuivende Ongelijkheid in Nederland*, edited by Jaap Dronkers and Wout C. Ultee, chapter 3, pp. 31–45. Assen: Van Gorcum.

De Jong, Uulkje, Jaap Dronkers, and Willem E. Saris. 1982. "Verandering in de Schoolloopbanen tussen 1965 en 1977: Ontwikkelingen in de Nederlandse Samenleving en in haar Onderwijs." *Mens en Maatschappij* 57:26–54.

DiPrete, Thomas A. and Markus Gangl. 2004. "Assessing Bias in the Estimation of Causal Effects: Rosenbaum Bounds on Matching Estimates and Instrumental Variables with Imperfect Instruments." *Sociological Methodology* 34:271–310.

Dronkers, Jaap. 1993. "Educational Reform in the Netherlands: Did It Change the Impact of Parental Occupation and Education?" *Sociology of Education* 66:262–277.

Dronkers, Jaap and Hans Bosma. 1990. “The Consequences of Educational Reform for the Unequal Educational Opportunities of Children.” *Zeitschrift für Sozialisationsforschung und Erziehunssoziologie* 10:268–278.

Drukker, David M. and Richard Gates. 2006. “Generating Halton Sequences using Mata.” *The Stata Journal* 6:214–228.

Duncan, Beverly. 1967. “Education and Social Background.” *The American Journal of Sociology* 72:363–372.

Erikson, Robert. 1984. “Social Class of Men, Women and Families.” *Sociology* 18:500–514.

Faasse, John H., Bart F.M. Bakker, Jaap Dronkers, and Huibert Schijf. 1987. “The Impact of Educational Reform: Empirical Evidence from Two Dutch Generations.” *Comparative Education* 23:261–277.

Fitzmaurice, Garrett M., Nan M. Laird, and James H. Ware. 2004. *Applied Longitudinal Analysis*. Hoboken, NJ: John Wiley & Sons., Inc.

Fox, John. 1997. *Applied Regression Analysis, Linear Models, and Related Methods*. Thousand Oaks: Sage.

Fox, John. 2000. *Nonparametric Simple Regression: Smoothing Scatterplots*. Quantitative Applications in the Social Sciences, 131. Thousand Oaks: Sage Publications.

Ganzeboom, Harry B.G. 1996. “Onderwijsexpansie en Onderwijskansen.” In *De Sociale Segmentatie van Nederland in 2015*, edited by Harry B.G. Ganzeboom and Wout C. Ultee, chapter 2, pp. 13–48. Den Haag: Wetenschappelijke Raad voor het Regeringsbeleid.

Ganzeboom, Harry B.G. and Paul M. De Graaf. 1989. “Intergenerationele Opleidingsmobiliteit in Nederland van Geboortecohorten 1891–1960.” *Sociale Wetenschappen* 32:263–278.

Ganzeboom, Harry B.G., Matthijs Kalmijn, and Jules Peschar. 1995. “Het Nederlandse Stratificatiepatroon in Internationaal Perspectief.” In *Verschuivende Ongelijkheid in Nederland*, edited by Jaap Dronkers and Wout C. Ultee, chapter 9, pp. 142–161. Assen: Van Gorcum.

Ganzeboom, Harry B.G. and Ruud Luijkx. 2004a. “More Recent Trends in Intergenerational Occupational Class Reproduction in the Netherlands 1970-2004, Evidence from an Expanded Database.” *Netherland Journal of Social Sciences* .

Ganzeboom, Harry B.G. and Ruud Luijkx. 2004b. "Recent Trends in Intergenerational Occupational Class Reproduction in the Netherlands 1970–1999." In *Social Mobility in Europe*, edited by Richard Breen, chapter 14, pp. 345–381. Oxford: Oxford University Press.

Ganzeboom, Harry B.G. and Donald J. Treiman. 2003. "Three Internationally Standardised Measures for Comparative Research on Occupational Status." In *Advances in Cross-National Comparison: A European Working Book for Demographic and Socio-Economic Variables*, edited by Jürgen H.P. Hoffmeyer-Zlotnik and Christof Wolf, chapter 9, pp. 159–193. New York: Kluwer Academic Press.

Ganzeboom, Harry B.G. and Donald J. Treiman. 2009. "International Stratification and Mobility File." `http://home.fsw.vu.nl/hbg.ganzeboom/ISMF/index.htm`, last accessed on 7/9/2009.

Gershuny, Jonathan, Michael Godwin, and Sally Jones. 1994. "The Domestic Labour Revolution: A Process of Lagged Adaptation." In *The Social and Political Economy of the Household: Social Change and Economic Life Initiative*, edited by Michael Anderson, Frank Bechhofer, and Jonathan Gershunny, pp. 151–197. Oxford: Oxford University Press.

Gesthuizen, Maurice, Paul M. De Graaf, and Gerbert Kraaykamp. 2005. "The Changing Family Background of the Low-Educated in the Netherlands: Socio-Demographic Resources." *European Sociological Review* 21:441–452.

Gill, Jeff. 2006. *Essential Mathematics for Political and Social Research*. New York: Cambridge University Press.

Goldthorpe, John H. 1983. "Women and Class Analysis: In Defence of the Conventional View." *Sociology* 17:465–488.

Greenstein, Theodore N. 2000. "Economic Dependence, Gender, and the Division of Labor in the Home: A Replication and Extension." *Journal of Marriage and the Family* 62:322–335.

Groot, Wim and Henriëtte Maassen van den Brink. 2000. "Overeducation in the Labor Market: A Meta-Analysis." *Economics of Education Review* 19:149–158.

Halton, John H. 1960. "On the Efficiency of Certain Quasi-Random Sequences of Points in Evaluating Multi-Dimensional Integrals." *Numerische Mathematik* 2:84–90.

Harrell, Frank E. 2001. *Regression Modeling Strategies: With Applications to Linear Models, Logistic Regression, and Survival Analysis*. New York: Springer.

Hartog, Joop. 2000. "Over-Education and Earnings: Where Are We, Where Should We Go?" *Economics of Education Review* 19:131–147.

Hauser, Robert M. and David L. Featherman. 1976. "Equality of Schooling: Trends and Prospects." *Sociology of Education* 49:99–120.

Hauser, Robert M. and Arthur S. Goldberger. 1971. "The Treatment of Unobservable Variables in Path Analysis." *Sociological Methodology* 3:81–117.

Heise, David R. 1972. "Employing Nominal Variables, Induced Variables, and Block Variables in Path Analysis." *Sociological Methods & Research* 1:147–173.

Holland, Paul. 1986. "Statistics and Causal Inference." *Journal of the American Statistical Association* 81:945–960.

Holm, Anders and Mads M. Jæger. 2008. "Selection Bias in Educational Transition Models: Theory and Empirical Evidence." Social Policy and Welfare Services Working Paper 11:2008, `http://www.sfi.dk/sw60623.asp`.

Hout, Michael and Thomas A. DiPrete. 2006. "What We Have Learned: RC 28's contributions to Knowledge about Social Stratification." *Research in Social Stratification and Mobility* 24:1–20.

Jeffreys, Harold. 1961. *Theory of Probability*. Oxford: Oxford University Press, 3rd edition.

Korupp, Sylvia E., Harry B.G. Ganzeboom, and Tanja Van Der Lippe. 2002. "Do Mothers Matter? A Comparison of Models of the Influence of Mothers' and Fathers' Educational and Occupational Status on Children's Educational Attainment." *Quality & Quantity* 36:17–42.

Korupp, Sylvia E., Karin Sanders, and Harry B.G. Ganzeboom. 2000. "Intergenerationele Overdracht van Status en Sekse-Typering van Beroepen: De Invloed van Vader en Moeder op hun Dochters en Zonen." *Mens en Maatschappij* 75:22–39.

Lauer, Charlotte. 2003. "Family background, cohort and education: A French-German comparison based on a multivariate ordered probit model of educational attainment." *Labour Economics* 10:231–251.

Little, Roderick J.A. and Donald B. Rubin. 2002. *Statistical Analysis with Missing Data*. Hoboken, NJ: John Wiley & Sons, Inc., 2nd edition.

Loader, Clive. 1999. *Local Regression and Likelihood*. New York: Springer-Verlag.

Loader, Catherine. 2005. `locfit`*: Local Regression, Likelihood and Density Estimation.* R package version 1.5-2, `http://www.locfit.info/`.

Long, J. Scott. 1997. *Regression Models for Categorical and Limited Dependent Variables*. Advanced Quantitative Techniques in the Social Sciences, 7. Thousand Oaks: Sage Publications.

Lucas, Samuel R. 2001. "Effectively Maintainted Inequality: Education Transitions, Track Mobility, and Social Background Effects." *American Journal of Sociology* 106:1642–1690.

Lucas, Samual R., Philip Fucella, and Mark Berends. 2007. "Neo-Classical Education Transitions: A Corrected Tale for Three Cohorts." Presentation at the International Sociological Association Research Committee Number 28 Meeting, Montréal Quebec, August 2007.

Maddala, Gangadharrao S. 1983. *Limited Dependent and Qualitative Variables in Econometrics*. Cambridge: Cambridge University Press.

Mare, Robert D. 1979. "Social Background Composition and Educational Growth." *Demography* 16:55–71.

Mare, Robert D. 1980. "Social Background and School Continuation Decisions." *Journal of the American Statistical Association* 75:295–305.

Mare, Robert D. 1981. "Change and Stability in Educational Stratification." *American Sociological Review* 46:72–87.

Mare, Robert D. 1993. "Educational Stratification on Observed and Unobserved Components of Family Background." In *Persistent Inequality: Changing Educational Attainment in Thirteen Countries*, edited by Yossi Shavit and Hans-Peter Blossfeld, chapter 15, pp. 351–376. Boulder: Westview Press.

Mare, Robert D. 1994. "Discrete-Time Bivariate Hazards with Unobserved Heterogeneity: A Partially Observed Contingency Table Approach." *Sociological Methodology* 24:341–383.

Mare, Robert D. 2006. "Response: Statistical Models of Educational Stratification — Hauser and Andrew's Models for School Transitions." *Sociological Methodology* 11:27–37.

McDonald, Gerald W. 1977. "Parental Identification by the Adolescent: A Social Power Approach." *Journal of Marriage and the Family* 39:705–719.

Meesters, Marion, Jaap Dronkers, and Huibert Schijf. 1983. "Veranderde Onderwijskansen? Een Derde Voorbeeld en Afrondende Conclusies." *Mens en Maatschappij* 58:5–27.

Mood, Carina. 2010. "Logistic regression: Why we cannot do what we think we can do, and what we can do about it." *European Sociological Review* 26:67–82.

Morgan, Stephen L. and Christopher Winship. 2007. *Counterfactuals and Causal Inference: Methods and Principles for Social Research*. New York: Cambridge University Press.

Neuhaus, John M. and Nicholas P. Jewell. 1993. "A Geometric Approach to Assess Bias Due to Omited Covariates in Generalized Linear Models." *Biometrika* 80:807–815.

Neuhaus, John M., John D. Kalbfleisch, and Walter W. Hauck. 1991. "A Comparison of Cluster-Specific and Population-Averaged Approaches for Analyzing Correlated Binary Data." *International Statistical Review* 59:25–35.

Nieuwbeerta, Paul and Susanne Rijken. 1996. "Educational Expansion and Educational Reproduction in Eastern Europe, 1940–1979." *Czech Sociological Review* 4:187–210.

O'Rand, Angela M. and John C. Henretta. 1982. "Delayed Career Entry, Industrial Pension Structure, and Early Retirement in a Cohort of Unmarried Women." *American Sociological Review* 47:365–373.

Peschar, Jules L. 1978. "Educational Opportunity Within and Between Holland and Sweden: The Semi-Experimental Approach." *Sociologische Gids* 25:273–296.

Peschar, Jules L. 1987. *Zo Vader — Zo Zoon, Zo Moeder — Zo Dochter? Vergelijkende Analyses naar de Processen van Statusverwerving en Onderwijsmobiliteit in Nederland, Hongarije en Polen*. Lisse: Swets & Zeitlinger.

Peschar, Jules L., Els Ten Vergert, and Roel Popping. 1986. "From Father to Son and From Father to Daughter: Educational Mobility in Hungary and the Netherlands for the Birth Cohorts from 1925 to 1955." *Quality & Quantity* 20:377–403.

R Development Core Team. 2005. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. `http://www.R-project.org`.

Raftery, Adrian E. 1995. "Bayesian Model Selection in Social Research." *Sociological Methodology* 25:111–163.

Rijken, Susanne. 1999. *Education Expansion and Status Attainment: A Cross-National and Over-Time Comparison*. Ph.D. thesis, ICS, Utrecht University, Utrecht.

Rijken, Susanne, Ineke Maas, and Harry B.G. Ganzeboom. 2007. "The Netherlands: Access to Higher Education — Institutional Arrangements and Inequality of Opportunity." In *Stratification in Higher Education*, edited by Yossi Shaviet, Richard Arum, and Adam Gamoran, pp. 266–293. Stanford: Stanford University Press.

Robert, Christian P. and George Casella. 2004. *Monte Carlo Statistical Methods*. Springer texts in statistics. New York: Springer-Verlag, 2nd edition.

Rosenbaum, Paul R. 2002. *Observational Studies*. New York: Springer, 2d edition.

Rosenbaum, Paul R. and Donald B. Rubin. 1983. "Assessing Sensitivity to an Unobserved Binary Covariate in an Observational Study with Binary Outcome." *Journal of the Royal Statistical Society. Series B* 45:212–218.

Royston, Patrick. 2004. "Multiple Imputation of Missing Values." *The Stata Journal* 4:227–241.

Royston, Patrick. 2005a. "Multiple Imputation of Missing Values: Update." *The Stata Journal* 5:188–201.

Royston, Patrick. 2005b. "Multiple Imputation of Missing Values: Update of `ice`." *The Stata Journal* 5:527–536.

Royston, Patrick. 2007. "Multiple Imputation of Missing Values: Further Update of `ice`, With an Emphasis on Interval Censoring." *The Stata Journal* 7:445–464.

Royston, Patrick. 2009. "Multiple Imputation of Missing Values: Further Update of `ice`, with an Emphasis on Categorical Variables." *The Stata Journal* 9:466–477.

Royston, Patrick and Mahesh K. B. Parmar. 2002. "Flexible Parametric Proportional-Hazards and Proportional-Odds Models for Censored Survival Data, With Application to Prognostic Modelling and Estimation of Treatment Effects." *Statistics in Medicine* 21:2175–2197.

Rumberger, Russell W. 1981. "The Rising Incidence of Overeducation in the US Labor Market." *Economics of Education Review* 1:293–314.

Shavit, Yossi and Hans-Peter Blossfeld. 1993. *Persistent Inequality: Changing Educational Attainment in Thirteen Countries*. Boulder: Westview Press.

Sieben, Inge, Johannes Huinink, and Paul M. De Graaf. 2001. “Family Background and Sibling Resemblance in Educational Attainment: Trends in the Former FRG, the Former GDR, and the Netherlands.” *European Sociological Review* 17:401–430.

Smith, Herbert L. and Paul P.L. Cheung. 1986. “Trends in the Effects of Family Background on Educational Attainment in the Philippines.” *The American Journal of Sociology* 91:1387–1408.

StataCorp. 2007. *Stata Statistical Software: Release 10*. College Station, TX: StataCorp LP.

Statistics Netherlands. 2008. “Statline.” `http://statline.cbs.nl/`, accessed on June 12, 2008.

Train, Kenneth E. 2003. *Discrete Choice Methods with Simulation*. Cambridge: Cambridge University Press.

Tutz, Gerhard. 1991. “Sequential Models in Categorical Regression.” *Computational Statistics & Data Analysis* 11:275–295.

UNESCO. 1997. *ISCED 1997, International Standard Classification of Education*. UNESCO.

Van der Ploeg, Sjerp. 1994. “Educational Expansion and Returns on Credentials.” *European Sociological Review* 10:63–78.

Vrooman, Cok and Jaap Dronkers. 1986. “Changing Educational Attainment Process: Some Evidence From the Netherlands.” *Sociology of Education* 59:69–78.

Wolbers, Maarten H.J. 1998. *Diploma Inflatie en Verdringing op de Arbeidsmarkt: Een Studie naar Ontwikkelingen in de Opbrengsten van Diploma’s in Nederland*. Ph.D. thesis, ICS, University of Nijmegen.

Wolbers, Maarten H.J. and Paul M. de Graaf. 1996. “Onderwijsexpansie of Meritocratisering? Een Verklaring voor de Afgenomen Samenhang tussen Ouderlijk Sociaal Milieu en het Hoogst Bereikte Opleidings Niveau van Kinderen.” *Tijdschrift voor Onderwijsresearch* 21:117–132.

Wolbers, Maarten H.J., Paul M. De Graaf, and Wout C. Ultee. 2001. “Trends in Occupational Returns to Educational Credentials in the Dutch Labor Market: Changes in Structures and in the Association?” *Acta Sociologica* 44:5–19.

Yamaguchi, Kazuo. 2002. “Regression Models with Parametrically Weighted Explanatory Variables.” *Sociological Methodology* 32:219–245.

Faculty of Social Sciences, VU University Amsterdam

Dissertation series

- M.M. Meijer, (2004), *Does Success Breed Success? Effects of News and Advertising on Corporate Reputation*. Aksant Academic Publishers: Amsterdam, isbn 90 5260 145 3.
- J.J. Woldendorp, (2005), *The Polder Model: From Disease to Miracle? Dutch Neo-corporatism 1965-2000*. Amsterdam, isbn 90 77472 05 3.
- Y.P. Ongena, (2005), *Interviewer and Respondent. Interaction in Survey Interviews*, Yfke Ongena, Amsterdam, isbn 90 9020070 3.
- M.E.D. Lamboo, (2005), *Integriteitsbeleid van de Nederlandse Politie*. Eburon: Delft, isbn 90 5972 088 L.
- T.E. Aalberts, (2006), *Politics of Sovereignty*. Vrije Universiteit Amsterdam, isbn 10: 90 9020656 6 / isbn 13: 978 90 9020656 1.
- E. Zwart, (2006), *In Pursuit of Comfort. The Transnationalisation Process of Malaysian Chinese Small and Medium Enterprises*. Vrije Universiteit Amsterdam, isbn 90 8659 024 1.
- H.Haenen, (2006), *Afrikaans Denken. Ontmoeting, Dialoog en Frictie. Een Filosofisch Onderzoek*. Buijten en Schipperheijn Motief, Amsterdam.
- A. Claver, (2006), *Commerce and Capital in Colonial Java. Trade Finance and Commercial Relations between Europeans and Chinese, 1820s-1942*. Vrije Universiteit Amsterdam.
- S. van der Pas, (2006), *Intergenerational Relationships of Older Adults. Family Structure, Contact and Norms*. Vrije Universiteit Amsterdam, isbn 10: 90 9022075 X / isbn13: 978 90 9021075 9.
- P. Mika, (2007), *Social Networks and the Semantic Web*. SIKS dissertation series No. 2007-03.
- K.E. Knibbe, (2007), *Faith in the Familiar. Continuity and Change in Religious Practices and Moral Orientations of the South of Limburg, The Netherlands*. Vrije Universiteit Amsterdam.
- C.L. Carabain, (2007), *Taking Too Much for Granted? A Study on the Measure of Social Attitudes*. Gildeprint Drukkerijen B.V., Enschede, isbn 978 90 8659 075 9.
- L.M. Cruz, (2007), *Repression and Rebellions in Southern Mexico. The Search for a Political Economy of Dignity*. Dutch University Press, Amsterdam, isbn 978 90 361 0050 2.
- C.E. de Vries, (2007), *European Integration and National Elections. The Impact of EU Issue Voting on National Electoral Politics*. Vrije Universiteit Amsterdam, isbn 13: 978 90 9022112 0.
- E. Kolthoff, (2007), *Ethics and New Public Management. Empirical Research into Effects of Businesslike Government on Ethics and Integrity*. Boom Juridische Uitgevers, isbn 978 90 5454 905 5.

- R. Vliegenthart, (2007), *Framing Immigration and Integration. Facts, Parliament, Media and Anti-Immigrant Party Support in the Netherlands*. Vrije Universiteit Amsterdam, isbn 978 90 9022279 0.
- L.L.J. van Noije, (2007), *The Democratic Deficit Closer to Home. Agenda Building Relations between Parliament and the Press, and the Impact of European Integration, in the United Kingdom, the Netherlands and France*. Vrije Universiteit Amsterdam, isbn 978-90-8659-161-9.
- R. de Man, (2007), *Where to Draw the Line? Antecedents and Consequences of the Integration of Work in Home Life*. Vrije Universiteit Amsterdam.
- F. Krijtenburg, (2007), *Cultural Ideologies of Peace and Conflict: a Socio-Cognitive Analysis of Giriyama Discourse (Kenya)*. Vrije Universiteit, isbn 978 90 8659 167 1.
- R. van Steden, (2007), *Privatizing Policing. Describing and Explaining the Growth of Private Security*. BJU Legal Publishers, isbn 978 90 5454 953 6.
- L. Reijn, (2008), *Identity at Work. Control and Commitment in Post Bureaucratic Organisations*. Vrije Universiteit Amsterdam, isbn 978 90 8659 180 0.
- B. Vis, (2008), *Biting the Bullet or Steering Clear? Politics of (Not-) Unpopular Welfare State Reform in Advanced Capitalist Democracies*. Vrije Universiteit Amsterdam, isbn 978 90 9022715 3.
- M.J.M. de Koning (2008), *Zoeken naar een 'Zuivere' Islam. Geloofsbeleving en Identiteitsvorming van Jonge Marokkaans-Nederlandse Moslims*. Bert Bakker.
- P. Wiepking (2008), *For the Love of Mankind. A Sociological Study on Charitable Giving*. Labor Grafimedia, isbn 978 90 77383 07 0.
- A. Supriyono (2008), *Buruh Pelabuhan Semarang. Pemogokan-Pemogokan Pada Zaman Colonial Belanda Revolusi dan Depublic 1900-1965*. Vrije Universiteit Amsterdam.
- G. van Noort (2008), *Bricks versus Clicks. A Relf-Regulation Perspective on Consumer Reponses in Online and Conventional Shopping Environments*. Vrije Universiteit Amsterdam, isbn 978 90 9022845 7.
- W. Stam (2008), *Does Networking Work? A Study of Social Capital and Performance among New Ventures in the Dutch Open Source Software Industry*. Vrije Universiteit Amsterdam, isbn 978 90 9023014 6.
- A. Wigger (2008), *Competition for Competitiveness: The Politics of the Transformation of the EU Competition Regime*. Rozenberg Publishers, isbn 978 90 9023150 1.
- Z. van der Wal (2008), *Value Solidity. Differences, Similarities and Conflicts between the Organizational Values of Government and Business*. Vrije Universiteit Amsterdam, isbn 978 90 9022991 1.
- H.C. van Vugt (2008), *Embodied Agents from a User's Perspective*. Vrije Universiteit Amsterdam, isbn 978 90 9023086 3.
- J. ter Horst (2008), *Weaving into Cambodia. Trade and Identity Politics in the (post)-Colonial Cambodian Silk Weaving Industry*. Vrije Universiteit Amsterdam, isbn 978 90 9023232 4.
- A. Berhe (2008), *A Political History of the Tigray People's Liberation Front (1975-1991): Revolt, Ideology and Mobilisation in Ethiopia*. Vrije Universiteit Amsterdam.

- H. van den Heuvel (2008), *Between Optimism and Opportunism. Deconstructing 'African Management' Discourse in South Africa*. Vrije Universiteit Amsterdam, isbn 978 90 9023118 1.
- K. Lasthuizen (2008), *Leading to Integrity. Empirical Research into the Effects of Leadership on Ethics and Integrity*. Vrije Universiteit Amsterdam, isbn 978 90 813277 1 8.
- I. Reci (2008), *Unemployed and Scarred for Life. Longitudinal Analyses of How Unemployment and Policy Changes Affect Re-employment Careers and Wages in the Netherlands, 1980-2000*. Vrije Universiteit Amsterdam, isbn 978 90 9023654 4.
- L.C. Horn (2009), *The Transformation of Corporate Governance Regulation in the European Union*. Vrije Universiteit Amsterdam, isbn 978 90 9023813 5
- B.E. Büscher (2009), *Struggles over Consensus, Anti-Politics and Marketing*. Vrije Universiteit Amsterdam.
- H.M. Koolma (2009), *Verhalen en Prestaties, een Onderzoek naar het Gedrag van Woningcorporaties*. Vrije Universiteit Amsterdam, isbn 978 90 8659 273 9
- J.J. van Wijk (2009), *Moving beyond Heroes and Winners*. Vrije Universiteit Amsterdam, isbn 978 90 8659 285 2
- H.A. Binnema (2009), *How Parties Change. EU Integration and the Flexible Response of Political Parties*. Vrije Universiteit Amsterdam, isbn 978 90 8659 309 5
- E. Grassiani (2009), *Morality and Normalcy in Asymmetrical Conflict: Distancing, Denial and Moral Numbing among Israeli Conscripts in Everyday Practices of Occupation*. Vrije Universiteit Amsterdam.
- G. van Koningsbruggen (2009), *Processing and Acceptance of Threatening Health Information: the Effects of Self-Affirmation*. Vrije Universiteit Amsterdam.
- D. van Bergen (2009), *Suicidal Behaviour of Migrant Women in the Netherlands*. Vrije Universiteit Amsterdam.
- J. Perry (2009), *Goodwill Hunting: Accounting and the Global Regulation of Economic Ideas*. Vrije Universiteit Amsterdam.
- A. Linden (2009), *Besmet. Levenslopen en Motieven van Extreem-Rechtse Activisten in Nederland*. Vrije Universiteit Amsterdam.
- A. Schakel (2009). *A Postfunctionalist Theory of Regional Government. An Inquiry into Regional Authority and Regional Policy Provision*. Vrije Universiteit Amsterdam.
- I. Plaisier (2009). *Work and Mental Health: Studies on the Impact Characteristics, Social Roles and Gender*. Vrije Universiteit Amsterdam, isbn 987 90 86593903
- J. Roeland (2009). *Selfation: Dutch Evangelical Youth between Subjectivization and Subjection*. Vrije Universiteit Amsterdam.
- A. Vliegenthart (2009), *Transnational Actors and Corporate Governance Regulation in Postsocialist Europe*. Vrije Universiteit Amsterdam.
- S. Bhagavatula (2009). *Weaving Social Networks. Performance of Small Rural Firms in India as an Outcome of Entrepreneurs' Social and Human Capital*. Vrije Universiteit Amsterdam.
- A. van Essen (2009). *Seeking a Balance: The Emergence of New Public Management in Health Care Regulation*. Vrije Universiteit Amsterdam, isbn 978 90 8659349 1

- J. Soons (2009). *Love Life and Happiness: A Study of Partner Relationships and Well-being in your Adulthood*. Vrije Universiteit Amsterdam, isbn 978 90 6984 596 8
- J. de Raadt (2009). *Contestable Constitutions: the Politics of Institutional Conflict in Post-Communist East Central Europe*. Vrije Universiteit Amsterdam, isbn 978 90 86593965
- B. Tasew (2009). *Methaphors of Peace and Violence in the Folklore Discourses of South Western Ethiopia: A Comparative Study*. Vrije Universiteit Amsterdam.
- D. Gberie (2010). *From Chains to Neckties: Liberated Africans and the Making of Modern Sierra Leone*. Vrije Universiteit Amsterdam.
- M. Guiaux (2010). *Social adjustment to widowhood: Changes in personal relationships and loneliness before and after partner loss*. Vrije Universiteit Amsterdam.
- W. Reedijk (2010). *Effecten van religieuze dialogen tussen Joden, Christenen en Moslims op eigen identiteit*. Vrije Universiteit Amsterdam.
- M. Hellendoorn (2010). *Opkomst en Neergang van het Kostwinnaarschapsysteem*. Vrije Universiteit Amsterdam, isbn 978 90 9025065 6
- E. Ersanilli (2010). *Comparing Integration. Host culture adoption and ethnic retention among Turkish immigrants and their descendents in France, Germany and the Netherlands*. Vrije Universiteit Amsterdam, isbn 978 90 5335 250 2
- T.S. Salverda (2010). *Sugar, Sea and Power: How Franco-Mauritians balance continuity and creeping decline of their elite position.* Vrije Universiteit Amsterdam.
- G.A. Moerman (2010). *Effecten van doorvraagtactieken in interviews met open vragen*. Vrije Universiteit Amsterdam, isbn 9789086594436

Expected

- H.L. Duijnhoven (2010). *For Security Reasons. Narrative about Security practices and organizational change in the Dutch and Spanish railway sector*. Vrije Universiteit Amsterdam, isbn 9789086594382
- N. Tuan Anh (2010). *Kinship as Social Capital: Economic, Social and Cultural Dimensions of Changing Kinship Relations in a Northern Vietnamese Village*. Vrije Universteit Amsterdam, isbn 978 90 5335 271 7
- M.J. Kamsma (2010). *The Jewish diasporascape in the straits: An ethnographic study if Jewish Businesses across borders*. Vrije Universiteit Amsterdam, isbn 978 90 5335 260 1