

Using household panel data
to understand the intergenerational
transmission of poverty

**Stephen P. Jenkins and
Thomas Siedler, April 2007**



Institute for Social and Economic Research
University of Essex
Wivenhoe Park, Colchester, Essex CO4 3SQ, UK.

stephenj@essex.ac.uk
www.iser.essex.ac.uk

tsiedl@essex.ac.uk

CPRC Working Paper 74

Chronic Poverty Research Centre
ISBN: 1-904049-73-7

Abstract

This paper discusses how household panel surveys can be informative about the intergenerational transmission of poverty. We consider issues both of data and of the statistical methods that may be applied to those data. Although the data focus is on panel surveys from developed countries, we also briefly consider data availability in developing countries. We set out a list of survey data requirements for intergenerational analysis, and then discuss how the main household panel surveys in developed countries meet the criteria. In order to highlight the advantages and disadvantages of household panel surveys, the section also compares them with other types of longitudinal studies. Next, we review the estimation methods that have been used to examine the intergenerational transmission of poverty when using household panel surveys. Finally, we provide three examples of household panel surveys in developing countries (Indonesia, Malaysia and Mexico) that meet the data requirements for analysis of the intergenerational transmission of poverty.

Keywords: Poverty, intergenerational transmission, panel data, household surveys

Acknowledgements

This study was commissioned by the Chronic Poverty Research Centre (CPRC), through its 'Empirical Approaches to the Intergenerational Transmission of Poverty' theme, under the supportive direction of Kate Bird. We are grateful to the anonymous referees for extensive and detailed comments, not all of which were able to be addressed. The responsibility for the material and the views expressed in this paper lies with the authors alone, and not ODI, CPRC, or the University of Essex.

Professor Stephen Jenkins is Director of the Institute for Social and Economic Research (ISER) at the University of Essex.

Thomas Siedler is Senior Research Officer, ISER.

0 Executive Summary

- This paper discusses how household panel surveys can be informative about the intergenerational transmission of poverty. We consider issues both of data and of the statistical methods that may be applied to those data. Although the data focus is on panel surveys from developed countries, we also consider selected examples of data that are available for developing countries.
- We discuss five criteria concerning data requirements which may be used to assess the suitability of data for empirical analysis of intergenerational transmission of poverty issues. We refer to the following:
 - the availability of appropriate measures of well-being (and thence poverty status).
 - the availability of measures of other factors that are relevant to the intergenerational transmission process (e.g. parental education).
 - the ability to link data within families, most notably across generations, so that there is the fundamental information about outcomes for individual and about family background variables. There are also substantial advantages from having data about all siblings within a family in order to control for unobserved within-family factors that may affect outcomes.
 - The availability of a large sample that is representative of the target population, and that remains so over time. Maintenance of representativeness of a longitudinal data sets relates to survey design features such as the ‘following rule’ that prescribes which members of the base sample information is collected about at successive interview rounds, and also to issues such as minimizing sample drop-out (‘attrition’).
 - The availability of repeated observations on key variables such as income over a period of time. This facilitates longitudinal averaging of such variables to reduce the potential impact of measurement errors and transitory variation, and enables researchers to investigate issues such as whether the timing of poverty during childhood matters.
- We argue that household panel surveys can meet these data requirements relatively well, referring to examples of leading panels from developed countries. The advantages and disadvantages of household panel surveys are highlighted with discussion of how well other longitudinal survey designs can meet the five criteria. We refer to retrospective surveys, cohort panels, rotating panels, and linked data from administrative records.
- Given suitable longitudinal data, we argue that there are five main statistical approaches to identifying the key features of the process of the intergenerational transmission of poverty:
 - Parametric regression models with ‘level’ estimators;
 - Parametric regression models with ‘sibling difference’ estimators;
 - Parametric regression models with ‘instrumental variable’ estimators;
 - Non-parametric bounds estimators, and
 - Propensity score matching methods.
- An overview of each of these methods is provided. We emphasise that each approach requires a different set of assumptions, and no method is to be an overall ‘best buy’ applicable in most circumstances. Our view is that it is valuable to employ as many of the

methods as one can, as this provides a means to check the robustness of any conclusions drawn from the analysis.

- Although there is a growing number of household panel surveys in low-income and middle-income countries, many of these surveys are not suitable for intergenerational analysis of the type undertaken for developed countries. This does not mean that such analysis is entirely ruled out. We refer to three examples of three household panel surveys in developing countries (Indonesia, Malaysia and Mexico) which appear to meet most of the data requirements criteria cited earlier, and we also cite a number of other developing country panels.

Table of Contents

Abstract and Acknowledgments	i
0 Executive Summary.....	ii
1 Introduction.....	1
2 Survey Data Requirements	1
2.1 Household Panel Surveys	3
2.2 Retrospective Surveys.....	5
2.3 Cohort Panels.....	5
2.4 Rotating Panels	6
2.5 Record Linkage	6
3 Statistical Methods To Assess The Effect Of Growing Up Poor On Later Life Income ...	7
3.1 Parametric Regression Models: Level Estimates	7
3.2 Parametric Regression Models: Sibling Difference Estimates.....	9
3.3 Instrumental Variable Estimates.....	10
3.3 Non-Parametric Bounding Methods.....	11
3.4 Propensity Score Matching Methods.....	12
4 Household Panel Surveys In Developing Countries.....	13
5 Conclusions.....	15
6 References	16

List of tables

Table 1. Three developing country household panel surveys for intergenerational analysis	19
--	----

1 Introduction

This paper discusses how household panel surveys can be informative about the intergenerational transmission of poverty. We consider issues of both data and of the statistical methods that may be applied to those data. Although the data focus is on panel surveys from developed countries, we also consider data availability in developing countries. The paper should be read in conjunction with the complementary CPRC study by Behrman (2006).

The paper has four sections following this Introduction. Section 2 sets out a list of survey data requirements for intergenerational analysis, and then discusses how the main household panel surveys in developed countries meet these criteria. In order to highlight the advantages and disadvantages of household panel surveys, the section also compares them with other types of longitudinal studies. Section 3 discusses the estimation methods that have been used in previous research to examine the intergenerational transmission of poverty when using household panel surveys. Section 4 refers to three examples of household panel surveys in developing countries (Indonesian Family Life Survey, Malaysian Family Life Survey and Mexican Family Life Survey). We discuss how each of these could allow researchers to investigate the intergenerational transmission of poverty, and we also provide citations to some other developing country panel surveys.

2 Survey data requirements

What are the data requirements necessary for deriving estimates of intergenerational associations in poverty, and the transmission process? We discuss data requirements with reference to five criteria.

Measures of well-being. The data must include measures of well-being, e.g. income, wages, financial or physical capital, or non-monetary measures of welfare such as food consumption, nutrition, health or housing conditions. And measures are required to be observed for both children and parents (see below) – because intergenerational analysis is the explicit focus.

Other important measures. If one wishes to go beyond documenting simple intergenerational correlations in well-being or poverty, then one requires measures that summarize key aspects of the intergenerational transmission process (as described e.g. in our companion paper, Jenkins and Siedler 2007). Examples of these measures include information about parents' and children's age, education, health, housing conditions, family structure, employment histories and neighbourhood characteristics. There are additional variables that one may wish to have as well, in order to improve the quality of the estimates derived in a statistical sense. (We discuss these variables under the heading of 'instrumental variables' in Section 3.)

Family linkages. A basic requirement for intergenerational research is that one can successfully match data about parents with data about their children. To be clear: what one typically requires is not contemporaneous observations on some well-being measure for an individual and their parent(s), but observations on outcomes during adulthood for the 'child', and observations on various measures of family background, particularly those during the period when the 'child' was growing up. This is likely to refer to at least a decade or more prior to the current outcome for the 'child', and this long time interval places substantial constraints on the data collection process. For example, earnings observed at relatively young ages may not be a good measure of longer-run earnings, due to the unsettled nature of labour market careers at this life-cycle stage (Grawe 2006). Better measures of the extent to which poverty is transmitted from one generation to the next are likely to require earnings

information from parents and children in the middle of their working careers (see e.g. Haider and Solon 2006). Another problem is how to adequately measure income or earnings because (1) not all surveys in developing countries collect information on income or consumption expenditures (Montgomery *et al.* 2000); and (2) a large part of the population has erratic employment, works in self-subsistence agriculture, or are employed in the informal sector.

Another sort of family linkage that is desirable from the point of view of improving the quality of the estimates derived is information about siblings from the same family. As we explain in Section 3, information about siblings helps researchers to control for unobservable family effects that might have an influence both on poverty during childhood and poverty later in life.

Large and representative samples. Research about intergenerational transmission requires large and representative samples, just as any other type of statistical analysis does. Small samples constrain the extent to which one can undertake breakdowns in analysis of different population subgroups and, more generally, the smaller the sample, the more likely that estimates are more prone to sampling variability and hence less reliable. Non-representative samples potentially limit the extent to which findings can be generalised. (An important prior issue is, of course, the decision concerning what the population of primary interest is.) Non-representativeness may arise through non-response by respondents, who either do not respond at all, or who do not respond to particular questions. Income data are often said to be particularly susceptible to this.

Ensuring the on-going representativeness of a longitudinal survey is a particular issue. On the one hand there is the desire to minimize the prevalence of sample drop-out ('attrition'), and non-random drop-out in particular. Another issue is how the survey design maintains representativeness of its target population over time, in particular for representing new entrants over time, persons and families, into that population (Buck *et al.* 1996). These design features of a longitudinal survey as usually labelled as 'following rules'. We would emphasize that the unit that is followed over time is the individual person, not the household or family. The reason for this is that it is impossible to define a longitudinal family or household in a rigorous way, because households fuse and split over time. (Individuals' income and poverty status are usually measured using information about the household within which they live, but that measurement issue is distinct from the issue of the unit that is tracked over time.)

Non-representativeness is a different issue from the coverage of the survey. If one's interest is only in a particular region of a country (say) and only this area is sampled, this may provide data that are representative of the target population. For example, a large random sample of India's most populous state, Uttar Pradesh, with more than 175 million inhabitants, could provide valuable data for researchers interested in studying the intergenerational persistence of poverty in that region. Whether the findings could be generalized to India as a whole is a different matter, and would depend on the extent to which mobility processes differed across regions. There are also issues of interpretation for a region-specific longitudinal survey that would arise if the survey did not 'follow' original sample members and re-interview if they left the sample region, e.g. to look for work elsewhere.

Repeated observations over a period. Income and other measures of well-being are essential for any study of poverty, but there is often concern expressed that income and related monetary measures may be subject to measurement errors or transitory fluctuations. If this is the case, then an observation on income at one point in time (whether for the parental generation or the child's generation) may be unreliable. In general, the more waves of data that are available for each generation, the better it is for estimating intergenerational transmission processes: averaging income over a number of years smoothes out and thereby minimize measurement error problems (Behrman and Taubman 1990; Solon 1992).

Another reason for having repeated observations, at least over childhood, is that it enables researchers to investigate whether the timing of poverty matters – whether e.g. poverty or low income more generally has a greater impact in early childhood years relative to later ones. If surveys provide asset information, it may be possible to use this information as a proxy for longer-run household income measures (see e.g. Behrman and Knowles 1999).

2.1 Household panel surveys

Household panel surveys in OECD countries meet most of the criteria that we have mentioned. We refer in particular to the most well-known and most widely-used long-running surveys such as the US Panel Study of Income Dynamics (PSID, which began in 1968), the German Socio-Economic Panel (GSOEP, 1984), and the British Household Panel Survey (BHPS, 1991). There are many other household panel surveys in OECD countries, e.g. in Sweden, the Netherlands, the European Community Household Panel and, more recently, in Canada, Switzerland, Australia and New Zealand, but many of them are not so well suited for intergenerational poverty analysis. For the younger panel surveys, this is particularly because the window of observation does not yet span two generations sufficiently well (hence failing the key family linkage criteria).

Household panels have a design according to which most variables are collected prospectively, as the survey proceeds from wave to wave. The advantage of that is detailed information may be collected about income, and related variables (e.g. various consumption expenditures and items, durables). For the ‘child’ generation, there are typically repeated observations over several key variables such as income. Contemporaneous collection of information assists reliability because respondents do not have recall very detailed pieces of information (the longer the recall interval, the more likely that memory may be systematically at error). We return to the recall bias issue shortly.

The prospective design means that information about two generations is not straightforwardly available, except with the passage of time. For example, most of the intergenerational analysis using US PSID data began only once the survey had been running for more than two decades, by which time the children of the original respondents had become respondent adults themselves.

The alternative means of undertaking intergenerational panels to date – primarily applied to the BHPS and the GSOEP – is to exploit the fact that these panels also included extensive retrospective recall question modules soon after the panels began that refer to periods of time well before the panel itself began. The retrospective life history information is used to provide information about family background, and the panel itself is used to provide information about later-life outcomes. The method is grounded on the fact that children become panel respondents in their own right at around age 16 and are then followed over time. By construction, these children lived with a parent who was (and may well remain) a panel member too, and so one can match the retrospective recall data for a parent with the panel data on outcomes and other variables for ‘children’. This method can provide relatively large samples for intergenerational analysis. One potential disadvantage is that the family background measures relating to family income during childhood, derived from the retrospective recall data, often do not refer to income specifically, but related variables. For example, Ermisch *et al.* (2001) using the BHPS recorded someone as experiencing poverty during childhood if both parents were not in paid work for at least one month in any one of the first sixteen years of life of the person concerned. (Information about parental work was derived by retrospective recall.) Some childhood income data may be available, but it may refer to only one year during childhood, which may be contrary to the ‘repeated observations’ criterion. The number of observations (and thence sample size) for which income during childhood may be available depends partly on how mature the panel is – the longer-running the panel, the more there are. In the limit (as with the PSID), income is available covering the

whole of childhood, and derived entirely from the panel (without matching in retrospective information).

We discuss usefulness of retrospective information in combination with current information further in Section 4, with reference to the potential of the Indonesia Family Life Survey for estimating intergenerational association in poverty in Indonesia.

A particular advantage of household panel survey designs is that, by construction, they collect information about all individuals within a given household at the time of the interview, including all siblings. And all children of a respondent household eventually become panel members in their own right (more on this below). The observation of multiple siblings meets another data desideratum.

With respect to representativeness, we focus on aspects that are particularly relevant to longitudinal surveys. First, there is the following rule by which household panel surveys maintain the on-going cross-sectional representativeness of the (non-immigrant) population, which is as follows. Define the adults and children in the representative sample of households in the first wave as 'original sample members' (OSMs). (Observe that the co-resident children sampled at wave one are not necessarily all the children the parents ever had, as some may have already left home or died. Similarly the adults present at wave one may not include both birth parents of a given child, e.g. of earlier parental divorce or death.) In subsequent waves, attempt interviews with all adult members of all households containing either an OSM or an individual born to an OSM whether or not they were members of the original sample, and regardless of whether the individual lives in the same household or residence as at the previous interview.

This following rule underlies the design of virtually all household panels in industrialised countries. However, differences exist with respect to treatment of new panel members who subsequently stop living with an OSM. In most surveys in developed countries, including the PSID and the BHPS, these people are not interviewed again (unless they retain an important relationship with a sample member, such as parent). By contrast, the GSOEP has, since wave 7, followed and interviewed all panel members, regardless of their relationship to the OSM. It is important to note that this sort of tracking rule is used in only a few developing country surveys. (The three examples of household panel surveys from developing countries discussed in Section 4 do apply this type of tracking rule.)

We emphasize the nature of the following rule partly in order to stress how rules that are not similar to the one just outlined may lead to problems for analysis. In particular, we refer to panel surveys that are residence-based, and interviewers returning to particular addresses rather than particular people. If a household has split up because of divorce, or children have left the household to live elsewhere, then the group of people who is interviewed (those remaining) is potentially a non-random sample, and so may lead to biased estimates. Rosenzweig (2003) reports considerable biases in the Bangladesh Nutrition Survey because the survey re-interviewed only those individuals that were still living in the household originally surveyed.

The other dimension of non-response that is particularly important for household panel surveys – indeed for all panel surveys – is selective sample drop-out (attrition). Attrition is a problem that potentially increases the longer the panel is, and hence is a feature that conflicts with the distinct advantages of having longer panels that we have already discussed. Attrition reduces sample size, and also introduces potential non-representativeness if sample drop-out is non-random. The latter case occurs when some individuals are systematically more likely to drop out of the panel than others. For example, if poorer respondents are more likely to drop out of the study, estimates of the degree of intergenerational transmission of poverty may be biased. Non-random panel attrition and

non-representativeness are issues that are often mentioned but not always addressed, at least in the context of intergenerational research. However, see Fitzgerald, Gottschalk, and Moffitt (1999) for an important exception.

In order to highlight the advantages and disadvantages of household panel surveys relative to the data requirement check-list, we now briefly compare their main features with other types of longitudinal survey designs. For a more extended discussion, see Buck, Ermisch, and Jenkins (1996).

2.2 *Retrospective surveys*

In retrospective surveys, individuals are typically interviewed only once and they provide retrospective information using recall. The advantages of retrospective surveys are simplicity, cheapness (mainly because there is only a single interview; respondents do not have to be tracked over time and place, etc.), together with the immediate availability of longitudinal information (since one does not have to wait for a second interview to measure change). This might allow researchers to investigate various aspects of intergenerational transmission of poverty with only one cross-sectional sample. Relatively large sample sizes of retrospective surveys allow researchers analysis of subjects who might be few in relative numbers in the population, such as lone parents or minority groups.

The principal disadvantage of relying on recall data is that the information about the past is dependent on respondents' recall of events, and the accuracy of this is questionable for many of the measures that are of interest. People are unlikely to remember very well earnings or income levels beyond the immediate past, or may do so with error, and these are fundamental to intergenerational studies of poverty. Major lifetime events such as getting married or divorced, or having a child, and their dates of occurrence are more likely to be remembered with reasonable accuracy. Measures of parental background of the sort that we seek in the current context are likely to be of relatively poor quality compared to those available from panel surveys as described earlier, and so retrospective surveys have been little used for intergenerational analysis. (The major exception to this is the analysis of social class mobility, where social class is based on information about employment and employment relations, collected by retrospective recall for the parental generation. See e.g. Goldthorpe 1980.) The very nature of recall data makes it more difficult to get multiple sibling data of the same quality as from a household panel survey, or to get repeated observations over a period on variables such as income that are of satisfactory quality.

2.3 *Cohort panels*

A cohort survey, a longitudinal survey focussing on the individuals from a specific birth cohort of the population (or some subsample of this), is the simplest example of a single indefinite life panel. By construction, the definition of membership of this group cannot change over time and so the 'following rule' for a cohort panel is simple: attempt interviews with all original sample members. Information might be collected at each interview about other persons in a sample member's household or about parents, but no attempt is made to follow these people: they cannot become sample members in their own right. Examples of cohort panels include the UK National Child Development Study (NCDS) following all children born in particular week in 1958, the British Cohort Study (BCS70) following all children born in particular week in 1970, and the Millennium Cohort Survey (MCS), following a sample of children born in 2000/1. In each case, data collection is not undertaken annually, as with household panel surveys, but at intervals, e.g. in the NCDS at ages 0, 7, 11, 16, 23, 33, and 42.

The infrequency of interviews mean that the chances of collecting repeated observations on income, or the possibility of examining income timing aspects, are small. Moreover because

most of these surveys started life primarily as medical studies, the amount of information collected about the socioeconomic aspects of family background at young ages is not as extensive as for e.g. the major household panels. (The medical and developmental information that is collected is, of course, something that is not collected by household panels in such extensive detail.) By design, information is collected about only one family member, and not about all siblings. Like household panels with prospective data collection, the potential for intergenerational analysis of poverty requires the panel to mature so that observations on both the family background and outcomes during adulthood can be collected. Attrition is also an issue (as with household panels).

The NCDS and BCS70 are examples of cohort panels that have matured sufficiently long to be used extensively for intergenerational analysis in general and, more recently, the intergenerational transmission of poverty in particular. Blanden and Gibbons (2006), for example, used the NCDS and BCS70 to examine the transmission process and how it had changed over time. Teenage poverty status was measured using data on the family income reported by the parents when the child was aged 16. Data from the interviews with respondents when they were their 30s (and also at the age of 42 in the NCDS) provided measures of poverty during adulthood.

2.4 Rotating panels

A rotating panel survey consists of a succession of separate panel surveys with staggered starting times. An initial sample of individuals is selected and interviewed a pre-determined number of times, often at intervals shorter than a year. During the life of this first panel, a new sample is selected, followed, and interviewed in the same way as the first. Subsequent panels are constructed similarly. As a result, respondents are being continuously rotated out of the panel and the number of panel participants replenished by those being rotated into the survey. The US Survey of Income and Program Participation is an example of such a survey. Since individuals are interviewed for a fixed number of times only, and the length of time spanned by these interviews is relatively short, information on income and earnings spanning two generations is limited (family socioeconomic background data is limited). This virtually rules out the use of rotating panels for analysing intergenerational transmission of poverty. Rotating panel designs are used by many panels in developing countries. Examples are the National Urban Employment Survey in Mexico, the Labour Force Survey in South Africa, and the Brazilian Monthly Employment Survey.

2.5 Record linkage

Another way of creating longitudinal data is through linkage of personal records from existing temporally-separate data sources. In this way, longitudinal data can be collected without personal interviews. These data sets may be administrative records gathered for official purposes, e.g. of social security benefit administration records, income tax returns, or surveys such as national censuses. The Nordic countries are pioneers in establishment of very extensive intergenerational data sets created by linking together various registers. There is growing discussion of use of such data in other countries such as UK, but little action to date. We note that such data sets can only be created if suitable administrative record sources exist (and appropriately computerized) and this makes their use in most low-income countries much less likely. Another potential constraint in developing countries is that tax authorities may have very little information about incomes of rural households (e.g. farm and non-farm incomes). Also, administrative data do not cover individuals working in the informal economy, and this constitutes a considerable fraction of the economy in many developing countries (Pratap and Quintin 2006).

Longitudinal data created in this way have some advantages relative to household panel surveys. First, they usually have very much larger sample sizes. Hence, analysis can be

constructed for almost every population subgroup of interest, and sampling errors are reduced. Second, by not using interviews or respondent recall, there is no respondent burden or recall or reporting errors. However there are also some potential problems with using linked record data. First, linkage is often not possible, as a result of confidentiality or privacy restrictions relating to collection of the original data, or only on a non-representative basis. Second, analysis may be constrained by their being a smaller range of variables collected in comparison to household panel surveys. A third problem is that the linked record data set may only provide concurrent information, so that events or status during an intervening period (potentially several years long) may not be recorded.

For the types of linked record intergenerational data that has been compiled in the Nordic countries, most of the disadvantages cited are relatively minor. A recent study illustrating the data's potential for examining intergenerational mobility is that by Jäntti *et al.* (2006). The authors compare intergenerational mobility across six nations, using linked register data for Denmark, Finland, Norway, Sweden, together with survey data for the UK (NCDS) and USA (NLSY) with much smaller sample sizes.

3 Statistical methods to assess the effect of growing up poor on later life income

A range of econometric procedures is available when using household panel data, depending on the nature of the dependent variable and the time period over which observations are made. Each method provides a different approach to estimate the effect on later-life income and other attainment measures (e.g. earnings, income, health measures or education) of growing up in a low income family.

Experience of low income during childhood is not the only determinant on outcomes later in life; there are other factors that are both observed (e.g. parents' education) and unobserved (e.g. parental motivation and 'ability'). See our companion paper for further discussion (Jenkins and Siedler 2007). It is these other influences that we wish to control for in order to assess the 'true' causal effect of growing up poor. Because of the complexities of the processes leading to particular outcomes, there is no straightforward way in which to do this. There are different statistical techniques available, each of which utilises a different set of assumptions. If each method points to the same result, then researchers may claim some robustness for the overall conclusions; if results differ, then one learns which sorts of assumptions about the model are critical for drawing conclusions. We now describe and discuss five different methods available when using household panels:

1. Parametric regression models with 'level' estimators;
2. Parametric regression models with 'sibling difference' estimators;
3. Parametric regression with instrumental variable estimators;
4. Non-parametric bounds estimators, and
5. Propensity score matching methods.

See Ermisch *et al.* (2004), Wooldridge (2002), Manski (1995), and Caliendo and Kopeinig (2007) for more detailed expositions than the non-technical exposition that follows.

3.1 Parametric regression models: level estimates

The prototypic empirical model of the process of intergenerational transmission of poverty can be described in linear regression form as follows:

Y_{ij}	<p>The outcome of interest for individual i from family j (e.g. poverty status)</p>	
$=$	βL_{ij}	<ul style="list-style-type: none"> • the family background variable of principal interest (e.g. poverty during childhood),
(1)	$+ \gamma X_{ij}$	<ul style="list-style-type: none"> • other observed influences (X_{ij}),
	$+ \alpha_j$	<ul style="list-style-type: none"> • an unobserved family-specific effect, and
	$+ v_{ij}$	<ul style="list-style-type: none"> • an unobserved individual-specific effect.

The outcome of interest Y_{ij} might measure whether a child is poor when an adult, or the number of years the child lived in a low-income family later in life, say. The key family background variable of primary interest, L_{ij} , might be a measure of whether the individual lived in a low income family during childhood. The parameter of most interest – the causal effect of living in a poor family during childhood – is the coefficient β . We are interested in not only its sign and magnitude, but also whether our estimate of it differs from zero in a statistically significant manner. Put another way, to be confident in asserting that growing up in poor family has a deleterious impact on attainment, one needs to find not only that the estimate of β is positive, but also that this did not arise by chance. Precision of the estimates can be assessed in the conventional way: the standard error associated with each estimated coefficient can be calculated, and the statistical significance can be calculated.

Equation (1) is a standard linear regression model specification, and it is well-known that application of standard estimation techniques lead to unbiased estimates of the model parameters, including β , as long as the observed influences on attainments (L_{ij} , X_{ij}) are uncorrelated with the unobserved influences (α_j and v_{ij}). If one makes this assumption, then the only remaining complication is that the outcome variable could be dichotomous rather than continuous, but this is straightforwardly addressed using standard probit or logit regression techniques. Note that instead of reporting β itself, one can report the ‘marginal effect’ instead, as the latter is in a metric that is more easily comparable across the different methods used. The marginal effect shows the change in the probability of achieving the (dichotomous) outcome variable that is associated with a one unit change in childhood poverty variable L (for example growing up in a poor family rather than in an affluent family). For examples of studies in this tradition, see our companion paper (Jenkins and Siedler 2007), and the references therein.

Application of this method provides researchers with a set of ‘level’ estimates (the reason for this label will become apparent shortly). Since many of the findings reported in the literature have been based on level estimates, they provide an important reference point for researchers.

Increasingly, however, it has been argued that it is implausible to assume that the observed and unobserved influences on attainment are uncorrelated with each other. The argument is related to assumptions about the nature of the attainment process. The family- or mother-specific influences on attainment that are not observed by the researcher (α_j) include factors such as maternal ‘ability’ or any other fixed effect that is common among siblings within the same family (for example motivation and work ethics). A proportion of these factors are likely to be inherited. The individual-specific family background influences on attainment are likely to depend not only on the given individual’s endowments of intelligence and ‘ability’ (for example), but also the endowments of that individual’s sibling(s): parental home investments in a child are likely to respond to that child’s capacity to benefit from it and also (differentially) to the capacities of their other children.

If the assumption that L_{ij} , α_j , and v_{ij} are uncorrelated is untenable, estimates of β using the 'level' method described so far will be subject to bias. The degree of bias can be shown to depend on several aspects of the intergenerational transmission process, including the degree of heritability of endowments, the extent to which parents reinforce or compensate for cross-children differences in their children's endowments, and the nature of the response of growing up poor to family- and child-specific factors (Ermisch and Francesconi 2001).

3.2 Parametric regression models: sibling difference estimates

How might one address these problems? Estimation of 'sibling difference' models is a common answer. The models are grounded on the observation that siblings share many family-specific characteristics that are relevant to the attainment process, for example 'biological or social parents, their parenting style, parents' social and cultural environments, housing and, to a large extent, neighbourhoods and schools' (Ermisch *et al.* 2004: 77). This means that one can control for the unobserved (and observed) effects that are common to siblings by looking at the differences in their attainment and relating these to sibling differences in the experience of childhood poverty.

If we use the symbol Δ to represent the difference between siblings, then we can rewrite equation (1) for a two child family as:

$\Delta Y_j = Y_{1j} - Y_{2j}$	The difference in attainments between siblings 1 and 2 is a function of
$=$	$\beta \Delta L$
(2)	<ul style="list-style-type: none"> • the sibling difference in family income or childhood poverty
	$+ \gamma \Delta X$
	<ul style="list-style-type: none"> • the sibling difference in other observed influences, and
	$+ \Delta v$
	<ul style="list-style-type: none"> • the sibling difference in individual-specific effects.

Estimation is based on observations of differences between pairs of siblings – hence the name of this model. (By contrast, the levels model discussed earlier was based on variables expressed in terms of levels of attainment rather than differences in such levels.) The key requirement for an unbiased estimator of β in the sibling difference model is that sibling differences in childhood poverty (ΔL) and sibling differences in unobserved individual effects (Δv) are uncorrelated, which is a weaker criterion for unbiasedness than the levels estimator required. Intuitively, the reason is that the family-specific effect (α_j) is eliminated when taking sibling differences – it does not appear in equation (2) – and so many of the contributions to bias in the levels estimator no longer play a role (those related to the degree of heritability and to family structure responses to family-specific factors).

The use of sibling difference models does not guarantee that estimates of β are unbiased, however. There remain some child-specific feedback factors contributing to potential bias that were also relevant to the estimation of the levels model. However, it can be shown that if the impact on childhood poverty of differences in children's individual endowments is negligible, the sibling difference estimate of β is unbiased (Ermisch *et al.* 2004: 77).

There is an additional complication that may bias sibling difference estimates (Ermisch and Francesconi 2001, Ermisch *et al.* 2004). Suppose a father develops a behavioural problem (unobserved by the researcher) – an example could be alcoholism – that does not affect the attainment prospects of his older child but does adversely affect the prospects of his younger child (because she is exposed to it for a longer time, say). Moreover, in addition, the problem precipitates family poverty. In this situation, the sibling difference estimator over-estimates the true effect of growing up poor: the estimate partly reflects the influence of the

unmeasured parental behaviour that is correlated with childhood poverty. Finally, it should be noted that if researcher aims to estimate the influence of parental poverty at a specific parental age (e.g. when the mother was aged 30–35) on children’s poverty later in life, then sibling estimators cannot be used because household income for the specified parental age range is the same for each sibling. With no variation between siblings, an estimate of the impact of age-specific parental poverty could not be derived using a sibling difference estimator.

In sum, the sibling difference approach is a useful addition to the modeller’s toolbox, but it does not come without costs. There remain assumptions that may not be satisfied. In addition, estimation is based only on families containing siblings. (Families with three or more siblings contribute more than one sibling pair.) This necessarily reduces sample sizes and so lowers the precision of estimates, other things being equal. A further complication is that the conditional logit version of the sibling difference model, appropriate for the case of a dichotomous outcome variable, uses even fewer observations (the sibling pairs for whom the outcome differs).

Also, the exclusion of individuals from one-child families means that one cannot derive an explicit estimate of β separately for this group: in effect, one must assume that the same process applies to them as for individuals with siblings. Another feature of sibling difference models is that estimates of the effects of any observed variable that has the same value for each sibling cannot be identified – these influences also get eliminated by the differencing procedure, even though they may be of substantial interest (This does not arise with level regressions.) An example is mother’s highest educational qualification (part of X).

3.3 Instrumental variable estimates

The method of instrumental variables (IV) provides another way of accounting for potential correlations between the explanatory variable of principal interest, such as whether the respondent experienced poverty during childhood, and the unobserved factors influencing the outcome variables under study. (Such correlations lead to bias in estimates of the coefficient β in equation (1).) The basic idea of the IV approach is to find an additional variable that determines childhood poverty status but which has no direct influence on the outcome variable. More specifically, one needs an observable variable (instrumental variable, denoted by Z_i), not included in equation (1), that satisfies two conditions. First, the instrument must not be correlated with the error term u_i , i.e. $cov(Z_i, u_i) = 0$, where u_i is the sum of the unobserved family-specific effect α_j and the unobserved individual-specific effect v_{ij} in equation (1). For notational convenience, we suppress the subscript j in the remainder of this section. If one has more than one instrument, one can test whether the first condition holds, but this is not possible if the number of instruments equals the number of potential endogenous variables (Sargan 1958).

The second condition is that the instrument has to be correlated with the potential endogenous variable (growing up poor), once all the other exogenous variables (X_i) have been netted out. This requirement can be written in terms of the following linear projection:

$$(4) \quad L_i = \delta_1 X_i + \delta_2 Z_i + \varepsilon_i,$$

with $E(\varepsilon_i) = 0$. Because the variables in the vector X_i and the instrument Z_i are assumed to be exogenous, the error term ε_i is uncorrelated with all variables on the right-hand side of equation (4) by construction. The second condition underpinning the validity of IV estimation is that the coefficient on Z_i differs from zero ($\delta_2 \neq 0$).

Both the endogenous variable L_i and the instrument Z_i may be either continuous or binary (Wooldridge, 2002). For instance, L_i could be equal to one if a person grew up in poverty,

and zero otherwise. Alternatively, L_i could represent the number of years or months a person lived in poverty during childhood or adolescence. If the two IV conditions hold, one can plug equation (4) into (1) and the new equation can be estimated by OLS or probit.

In practice, it may be difficult to find an instrument which satisfies the two IV conditions. The second condition can be tested, and recent work by Bound, Jaeger and Baker (1995), Staiger and Stock (1997) and Stock and Yogo (2005) point to potential problems with 'weak instruments'. The issue is that, if the partial correlation between the instrument Z_i and the endogenous variable L_i is weak, IV estimates may be biased as well. Finally, note that IV estimates tend to lead to less efficient coefficient estimates, and there is a risk that estimates are too inefficient to be informative (Murray 2006).

3.3 Non-parametric bounding methods

Another set of estimators is inspired by research on social experiments in which one investigates the effect of a given 'treatment' on an outcome. For example, growing up in poverty can be considered as the treatment variable, and one might study each of a number of dichotomous outcome variables, e.g. achieving a specific educational qualification level, or being in poor health, being unemployed or being poor as adult. What we are interested in is the 'treatment effect' of growing up poor on each outcome, i.e. the difference between the probability that a young adult would achieve the outcome if he grew up in a poor household (say) and the probability that the same young adult would achieve the outcome if, instead, he did not grow up poor.

The problem is that, with most survey data such as provided by household panel surveys, researchers do not know what the counterfactual outcomes are – what would have happened to the young adult who grew up poor if he had in fact grown up in a more affluent family or, similarly, what would have happened to the young adult who grew up in a non-poor family if he had grown up in a poor family.

Parametric regression models, combined with various assumptions about the intergenerational transmission mechanism, are one way of getting round this problem, as explained earlier. Alternatively, one might ask how much can be said about treatment effects in the absence of any parametric specification or assumptions about the transmission mechanism. Using the method proposed by Manski (1995), one can consider what bounds may be put on the treatment effect. The method works as follows (we draw on the detailed exposition of Ermisch *et al.* 2004).

Each of the two probabilities used to define the treatment effect (the probability of attaining the outcome if treated and the probability of attaining the outcome if not treated) can be written as the sum of two terms, each of which is the product of a conditional probability and an unconditional probability. Within each sum, one of the constituent conditional probabilities cannot be observed in survey data but, because they are probabilities, each of them must lie between zero and one. Substituting these extreme values for the unobserved conditional probabilities allows one to put an upper bound and a lower bound on the probability difference that defines the treatment effect.

These upper and lower bounds represent the limits between which the treatment effect – the non-parametric counterpart to β – must lie. In principle, the treatment effect may lie between -1 and 1 (because it is defined as a difference between two probabilities), i.e. of width 2. The bounds implied by Manski's method can be shown to have width 1, a substantial reduction. On the other hand, the bounds also include 0, and so the method does not put bounds on the sign of the treatment effect. (They do not allow us to say whether the impact of growing up in a poor family has a positive or negative effect on attainment, for example.) This is not very informative, especially if zero lies near the middle of the range defined by the bounds.

To tighten the bounds on the treatment effect, one can estimate them separately for groups of individuals with similar characteristics (Ermisch *et al.*, 2004; Pepper, 2000). For example, Jenkins *et al.* (2005) use 96 groups defined by age, year of birth, sex, and mother's highest educational attainment, mother's age at the child's birth and year of birth. This provides them with 96 sets of upper and lower bounds. Of particular interest, and what many researchers report, are estimates of the largest lower bound and the smallest upper bound (and their standard errors). The difference between these estimates is at most equal to one but may be smaller. This tightening of the bounds has the potential to provide a better indication of the magnitude of the treatment effect (Ermisch *et al.* 2004), as long as the relevant bounds are relatively precisely estimated (precision is related to the number of groups used and hence within-group sample numbers). Although Manski's bounds are helpful in narrowing down the range of the treatment effect, a potential problem with this approach is that by grouping the data into cells according to many observable characteristics the number of cells might become very large and the number of observations in each cell very small. This may decrease the statistical precision of estimates considerably. In addition, there might be important differences in observable characteristics which are continuous in nature – such as wages, income or number of years of education – which makes grouping data into cells even more complicated. Propensity score matching methods provide a convenient solution to this problem (Rosenbaum und Rubin 1983).

3.4 Propensity score matching methods

The aim of matching methods is to construct the missing counterfactual, i.e. to derive a control group or non-treated individuals who are very similar to treated individuals with respect to observable characteristics. A key assumption of propensity score matching methods is that characteristics that are unobservable to the researcher, such as individuals' ability or motivation or parents' parenting style, are independent of the treatment (i.e. independent of childhood poverty status in the current context), given observable characteristics X_i . Hence it is important that researchers have data to hand which contain a comprehensive set of socioeconomic characteristics in order to minimize any potential problems arising from unobserved variables.

When using propensity score matching methods, one starts by estimating the probability of having experienced poverty during childhood (likelihood of being in the treatment group) as a function of observable characteristics, both for those who grew up poor and those who did not. In a second step, propensity score matching compares socioeconomic outcomes for treated and non-treated individuals who have similar probabilities in having grown up poor. This can be done by estimating, in a first step, a (logit or probit) model regression model for the probability that each individual had experienced childhood poverty, conditional on observable characteristics of individual i from family j , for example child's health, parents' education, mother's age at birth or father's occupation:

$$(6) \quad \Pr(L_i = 1 \mid X_{ij}).$$

From the estimated parameters of equation (6), one can derive the predicted probability for each individual of experiencing childhood poverty, \hat{P}^L , and also of not experiencing it, \hat{P}^C . Then, one can find non-treated individuals comparable with each treated individual by looking for individuals from each group with the same predicted probability of the relevant event. Once individuals have been matched, one can estimate treatment effects by comparing, for example, the average outcome for the group that experienced childhood poverty with the average outcome for the matched comparable group who did not experience childhood poverty.

The matching can be done by applying one of a number of different algorithms suggested in the recent literature, for example nearest neighbour matching, radius matching, or kernel density matching (Caliendo and Kopeinig 2006). In large samples, all matching algorithms should yield in similar results (Smith 2000). For further discussion of matching methods, see Blundell and Costa Dias (2000), Caliendo and Kopeinig (2006), Heckman, Ichimura and Todd (1997), and Smith and Todd (2005).

One advantage of the propensity score matching methods is its flexibility, since no particular functional form needs to be specified in order to compute the treatment effect. In contrast, level estimations in equation (1) use the assumption of linearity in functional form to yields unbiased estimates.

4 Household panel surveys in developing countries

This section discusses the extent to which one can analyze the intergenerational transmission of poverty using the household panel survey data that are available for developing and transition countries.

To be sure, there are household panel surveys in many developing countries nowadays, for example, the Côte d'Ivoire Living Standard Survey, the Longitudinal Community Panel Database in Ecuador, the Ethiopian Rural Household Survey and the Indonesian Family Life Survey (IFLS), to name but a few. For detailed information about household panel studies around the world see the compilations provided in the *Keeping Track: a Guide to Longitudinal Resources* database at <http://www.iser.essex.ac.uk/ulsc/keeptrack/index.php>, the *Panel Studies around the World* database at <http://psidonline.isr.umich.edu/Guide/PanelStudies.aspx>, and the document summarizing 'Panel datasets in developing and transitional countries' compiled by David Lawson, Andy McKay and Karen Moore available at <http://www.chronicpoverty.org/pdfs/PanelDatasetsVersion1-July%202003.pdf>.

However, few household panel surveys in these countries allow the analysis of intergenerational transmission of poverty in the same way that one can use those in industrialised countries (cf. the list of data requirements in Section 2). First, sample sizes are often too small. Second, many panels in developing countries have only a few waves of data available (Baulch and Hoddinott 2000). For instance, the KwaZulu-Natal Income Dynamics Study in South Africa collected data for the years 1993 and 1998. Third, retrospective information on employment, marital status, earnings and income are not always collected. Fourth, some household panel studies in developing countries did not follow people once they moved out of the original household or dropped households once they split up. For instance, the first generation of the Indian Village Level Studies of the International Crop Research Institute for the Semi-Arid Tropics, which started in 1975 and ran for 10 years, dropped from the sample households that split (Rosenzweig 2003). However, the second generation of the village level studies which started in 2002 now also includes split-off households into the sample.¹ Recent studies by Thomas *et al.* (2001), and Foster and Rosenzweig (2002), suggest that panels that condition re-interviewing on residence provide non-random samples of the households in the population. One way of accounting for this is to jointly model the intergenerational transmission process and the sample selection implied by the sample design. Finally, we note that some panels cover specific geographic areas, or were designed for specific purposes (Baulch 2003; McCulloch and Baulch 2000). Whether conclusions derived from these data may be generalized to different regions or populations is an issue that needs to be considered – in the same way that it is arguable whether estimates of intergenerational mobility derived from US panel describe the situation in the UK as well.

¹ See http://www.icrisat.org/qt-mpi/Projects/prj2_link.htm for further information.

There have been major advances and improvements in collection and design of household panel studies in a number of developing countries (Grosch and Glewwe 1996; Rosenzweig 2003). To illustrate this, we focus on three examples of household panel surveys that provide or will provide useful information for studying intergenerational mobility in developing countries: the Indonesia Family Life Survey (IFLS), the Malaysian Family Life Survey (MFLS), and the Mexican Family Life Survey (MxFLS1). We focus our discussion on these three examples because they satisfy most of the data requirements for the intergenerational analysis previously discussed, and because providing a comprehensive survey of all potentially suitable longitudinal data sets available in developing countries is beyond the scope of this paper. We should stress that there exist developing country panels appropriate for studying intergenerational transmission of poverty in addition to the three we focus on. For example, the first wave of the Additional Rural Income and Rural Economic Development Surveys (ARIS-REDS) in 1971 is a representative sample of the rural population in India residing in one of the 17 major states. The survey started with 4,527 households in 259 villages. Re-interviews were conducted in 1982 and 1999. In 1999, split-off households from the original survey still residing in the same village were surveyed again (Foster and Rosenzweig 2004).

For a tabular summary, for each of the three panel studies considered, of the panel design, time periods, number of surveys, following rules, and other information relevant to estimation of intergenerational mobility, see Table 1 (at the end of the paper). The discussion here provides only a brief commentary.

The three panels have relatively large sample sizes. For instance, the first wave of the IFLS consists of more than 22,000 individuals in 7,224 households. Wave 1 of the MFLS consists of 1,262 households with an ever-married woman which were selected to be representative of Peninsular Malaysia in 1976. The MxFLS1 in 2000 consists of 8,400 households in 150 communities across Mexico, and the questionnaires follow the design of the IFLS, adapted to the Mexican context and further surveys are planned.

Both the IFLS and the MFLS were successful in minimising drop-out rates. In both the second and third wave of the IFLS, around 95 percent of households were successfully contacted again in the sense that at least one respondent from the household in the first wave was interviewed. Similarly, IFLS2 in 1997/98 succeeded in interviewing 91 percent of all panel respondents from IFLS1. This implies that attrition rates are comparable with those for most household panel studies in the United States and Europe.

Importantly for intergenerational research, one key advantage of the IFLS over many other panel studies in developing countries is that the survey follows respondents even if they have moved out of or split off from original households in IFLS1. This allows researchers to match children with their parents and grandparents even though they do not live in the same household. Another strength of the Indonesia Family Life Survey is its broad coverage of retrospective information on employment, marriage, fertility, migration and wages over the life course for all panel respondents aged 15 and above.

We discussed earlier how retrospective recall information may be combined with household panel survey data to derive intergenerational data. We referred to the example of the study by Ermisch *et al.* (2001) based on the BHPS, in which retrospective employment histories were used to construct a measure of childhood poverty experience for 'children'. In contrast to the BHPS, the IFLS surveys are not conducted every year. However, in addition to income/wage information in the year of the surveys, all three IFLS waves provide retrospective information about monthly net wages/salary/income for the past 5 years. Hence, for parents who participated in all three surveys and provided valid income/wage information, information for estimating intergenerational associations are available for the period 1988–2000. Similarly, for adult children who participated in IFLS2 one can extract

income information spanning the period 1992/93–1997/98. This enables researcher to investigate various aspects of intergenerational mobility in Indonesia.

Of course, the use of retrospective histories on income raises questions about the quality of the data (see the earlier discussion). Recalling income values accurately from five years ago is certainly more difficult than doing so for income or wages from last month or last year. Results based on variables with measurement error might bias the estimates. One possibility to minimise measurement error would be to compute averages in parents' incomes as suggested by Solon (1992). Furthermore, since the IFLS provides retrospective information not only on income, but also on type of employment, number of hours worked and occupation, researcher can use these additional data to validate retrospective income information. Finally, in order to evaluate the quality of retrospective information, IFLS2 panel respondents in 1997/98 were interviewed on employment history for a nine-year observation window. This allows researchers to compare data quality of IFLS1 and IFLS2 for the four-year period between 1989 and 1993 (Frankenberg and Thomas 2000).

5 Conclusions

We have considered how household panel surveys may be used to understand the intergenerational transmission of poverty from the perspective of both data requirements and the statistical methods that one might apply once one had the intergenerational data. We referred to five main data requirements, and argued that the leading household panel surveys in industrialised countries met these criteria reasonable well. Other longitudinal survey designs may, in principle, provide suitable intergenerational data – notably cohort panels and linked record register data – and have been used successfully to do so in a number of industrialised countries. In a developing country context, household panel surveys seem to provide the design of most use for intergenerational analysis at the present time.

Our discussion of statistical methods to investigate causal pathways in the intergenerational transmission process referred to four main methods, two regression-based procedures and two more non-parametric ones, all of which have been applied in earlier studies. Each of the different approaches has strengths and weaknesses. Since each relies on different assumptions or otherwise has different properties, no one method is likely to be an overall 'best buy' applicable in most circumstances. Our view is that it is valuable to employ as many of the methods as one can, as this provides a means to check the robustness of any conclusions drawn from the analysis.

Although there is a growing number of household panel surveys in low income and middle income countries, many of them are not suitable for intergenerational analysis of the sort more routinely undertaken in developed countries (see our companion paper, Jenkins and Siedler 2007, for a review of findings). However, this does definitely not mean that such analysis is entirely ruled out; quite the contrary. We have provided three examples of three household panel surveys in developing countries which appear to meet most of the data requirements discussed in Section 2, namely the Indonesian Family Life Survey, the Malaysian Family Life Survey and the Mexican Family Life Survey. These surveys provide sufficiently detailed information to allow researchers to shed light on the intergenerational transmission of poverty in these countries.

6 References

- Baulch, B. (2003), 'Moving out and into poverty – A Q2 Perspective', paper presented at the WB Workshop on "Moving Out of Poverty: Growth and Freedom from the Bottom-Up", World Bank, July 15–17, 2003.
- Baulch, B. and Hoddinott, J. (2000), 'Economic mobility and poverty dynamics in developing countries', *Journal of Development Studies*, 36(6), 1–24.
- Behrman, J. R. (2006), 'Methodological note: using micro data to understand better the intergenerational transmission of poverty in low income developing countries.' Working Paper 68, IDPM/Chronic Poverty Research Centre (CPRC), Manchester.
<http://www.chronicpoverty.org/resources/cp68.htm>.
- Behrman, J. R. and Knowles, J. C. (1999), 'Household income and child schooling in Vietnam', *World Bank Economic Review*, 13(2), 211–256.
- Behrman, J. R. and Taubman, P. (1990), 'The intergenerational correlation between children's adult earnings and their parents income: results from the Michigan Panel Survey of Income Dynamics', *Review of Income and Wealth*, 36(2), 115–127.
- Blanden, J. and Gibbons, S. (2006). *The Persistence of Poverty Across Generations*, The Policy Press, Bristol.
- Blundell, R. and Costa Dias, M. (2000), 'Evaluation methods for non-experimental data', *Fiscal Studies*, 21(4), 427–468.
- Bound, J., Jaeger, D. A. and Baker, R. M. (1995), 'Problems with instrumental variables estimation when the correlation between the instrument and the endogenous explanatory variable is weak', *Journal of the American Statistical Association*, 90(430), 443–450.
- Buck, N., Ermisch, J. and Jenkins, S. P. (1996), 'Choosing a longitudinal survey design: the issues', Occasional Paper 96-1, Colchester: University of Essex.
<http://www.iser.essex.ac.uk/pubs/occpapers/pdf/op96-1.pdf>
- Butz, W. P. and Da Vanzo, J. (1978), 'The Malaysian Family Life Survey', Summary report R-2351-AID. RAND, Santa Monica CA.
- Caliendo, M. and Kopeinig, S. (2007), 'Some practical guidance for the implementation of propensity score matching', *Journal of Economic Surveys*, forthcoming.
- Ermisch, J. and Francesconi, M. (2001), 'Family matters: impacts of family background on educational attainments', *Economica* 68(270), 1372–56.
- Ermisch, J., Francesconi, M., and Pevalin, D. J. (2001), 'Outcomes for children of poverty', Research Report No 158, Department for Work and Pensions, Corporate Document Services, Leeds.
- Ermisch, J., Francesconi, M., and Pevalin, D. J. (2004), 'Parental partnership and joblessness in childhood and their influence on young people's outcomes', *Journal of the Royal Statistical Society A*, 167(1), 69–101.
- Fitzgerald, J., Gottschalk, P., and Moffitt, R. (1999), 'The impact of attrition in the PSID on intergenerational analysis', *Journal of Human Resources*, 33(2), 300–344.

- Foster, A. D. and Rosenzweig, M. R. (2002), 'Household division and rural economic growth', *Review of Economic Studies*, 69(4), 839–70.
- Foster, A. D. and Rosenzweig, M. R. (2004), 'Agricultural productivity growth, rural economic diversity, and economic reforms: India, 1970-2000', *Economic Development and Cultural Change*, 52, 509–542.
- Frankenberg, E. and Thomas, D. (2000), 'The Indonesia Family Life Survey (IFLS): study design and results from waves 1 and 2', RAND, Santa Monica CA.
- Goldthorpe, J., in collaboration with Llewellyn, C and Payne, C. (1980). *Social Mobility and Class Structure in Modern Britain*, Clarendon Press, Oxford.
- Grawe, N. D. (2006), 'Lifecycle bias in estimates of intergenerational earnings persistence', *Labour Economics*, 13(5), 551–570.
- Grosh, M. E. and Glewwe, P. (1996), 'Household survey data from developing countries: progress and prospects', *American Economic Review, Papers and Proceedings*, 86(2), 15–19.
- Haider, S. and Solon, G. (2006), 'Life-cycle variation in the association between current and lifetime earnings', *American Economic Review*, 96, 1308–1320.
- Heckman, J., Ichimura, H. and Todd, P. (1997): 'Matching as an econometric evaluation estimator: evidence from evaluating a job training programme', *Review of Economic Studies*, 64(4), 605–54.
- Jäntti, M., Røed, K., Naylor, R., Björklund, A., Bratsberg, B., Raaum, O., Österbacka, E. and Eriksson, T. (2006), 'American exceptionalism in a new light: a comparison of intergenerational earnings mobility in the Nordic Countries, the United Kingdom and the United States', Discussion Paper 1938, IZA, Bonn.
- Jenkins, S. P., Francesconi, M. and Siedler, T. (2005), *The Impact of Family Structure During Childhood on Later-Life Attainment*. Report Series, Anglo German Foundation, London. <http://www.agf.org.uk/pubs/pdfs/1414web.pdf>.
- Jenkins, S. P. and Siedler, T. (2007), 'The intergenerational transmission of poverty in industrialized countries', Working Paper 75, IDPM/Chronic Poverty Research Centre (CPRC), Manchester. <http://www.chronicpoverty.org/resources/cp75.htm>.
- Manski, C. (1995), 'Nonparametric bounds on treatment effects', *American Economic Review, Papers and Proceedings*, 80, 25–37.
- McCulloch, N. and Baulch, B. (2000), 'Simulating the impact of policy upon chronic and transitory poverty in rural Pakistan', *Journal of Development Studies*, 36(6), 100–130.
- Montgomery, M. R., Gagnolati, M., Burke, K. A. and Paredes, E. (2000), 'Measuring living standards with proxy variables', *Demography*, 37(2), 155–174.
- Murray, M. P. (2006), 'Avoiding invalid instruments and coping with weak instruments', *Journal of Economic Perspectives*, 20(4), 111-132.
- Pepper, J. V. (2000), 'The intergenerational transmission of welfare receipt: a nonparametric bounds analysis', *Review of Economics and Statistics*, 83(3), 472–488.

- Pratap, S. and Quintin, E. (2006), 'The informal sector in developing countries', Research Paper 2006/13, World Institute for Development Economics Research, United Nations University, Helsinki.
- Rosenbaum, P. and Rubin, D. (1983), 'The central role of the propensity score in observational studies for causal effects', *Biometrika*, 70, 41–50.
- Rosenzweig, M. R. (2003), 'Payoffs from panels in low-income countries: economic development and economic mobility', *American Economic Review Papers and Proceedings*, 93(2), 112–117.
- Sargan, J. D. (1958), 'The estimation of economic relationships using instrumental variables', *Econometrica*, 26, 393–415.
- Smith, J. (2000), 'A critical survey of empirical methods for evaluating active labor market policies', *Schweizerische Zeitschrift für Volkswirtschaft und Statistik*, 136(3), 1–22.
- Smith, J. and Todd, P. (2005), 'Does matching overcome LaLonde's critique of nonexperimental estimators?', *Journal of Econometrics*, 125(1–2), 305–353.
- Solon, G. (1992), 'Intergenerational income mobility in the United States', *American Economic Review*, 82(3), 393–408.
- Stock, J. H. and Yogo, M. (2005), 'Testing for weak instruments in IV regressions', in Andrews, D. and Stock, J. H. (eds.) *Identification and Inference for Econometric models: a Festschrift in honor of Thomas Rothenberg*, Cambridge University Press, Cambridge, 80–108.
- Thomas, D., Frankenberg, E. and Smith, J. P. (2001), 'Lost but not forgotten: attrition and follow-up in the Indonesian Life Survey', *Journal of Human Resources*, 36(3), 556–592.
- Staiger, D. and Stock, J. (1997), 'Instrumental variables regression with weak instruments', *Econometrica*, 65(3), 557–586.
- Wooldridge, J. M. (2002), *Econometric Analysis of Cross-Section and Panel Data*, MIT Press, Cambridge MA.

Table 1. Three developing country household panel surveys for intergenerational analysis

Household panel survey	Time period / Number of survey years / Sample size	Interview mode	Following rule	Retrospective information	Intergenerational information
Indonesia Family Life Survey (IFLS)	1993/94, 1997/98, 2000; 3 surveys. The first wave covers a sample of 7,224 households (more than 22,000 individuals) in 13 provinces in Indonesia, representing around 83 percent of the Indonesian population.	Interviews in 1993/94 (IFLS1) were conducted with (a) household head and spouse; (b) two randomly selected children of the head and spouse aged 0 to 14 (interviewed by proxy); (c) household member aged 50 and above and spouse, randomly selected from remaining members; (d) for a randomly selected 25 percent of the households, an individual aged 15 to 49 and their spouse, randomly selected from remaining members. Children in category (b) include biological, stepchildren or adopted children of the household head or spouse as well as any children fostered to any adult in the household. A household was defined as a group of people whose members reside in the same dwelling and share food from the same cooking pot. Maximum number of interviewed adult household members is 4. Proxy interviews are conducted for children, infants and temporarily absent household members.	Tracking and interviewing individuals who had moved or split off from origin IFLS1 households. In 1997/98 (IFLS2), 92% of IFLS1 respondents who provided detailed individual-level information or had been aged 26 or older in IFLS1 were re-interviewed. In origin households investigators sought to interview all members. In split-off households the aim was to interview original household members, their spouses and any of their biological children living in the new household. In IFLS2, children aged 11–14 were allowed to give their own interviews if the felt comfortable in doing so.	Employment, marriage, fertility, migration and wages over the life course for all panel respondents aged 15 and above.	The survey collects detailed information on economic status of individuals and households, including consumption, expenditures, earnings, non-labour income and wealth, and inventory of household consumption. Household questionnaire contains information about revenue, expenses, and value of assets of household-owned agricultural and non-agricultural businesses. Respondent who answers household questionnaire provides information on labour income for all individuals age 10 and above in the household who were not selected for detail interview, as well as household-level aggregate amounts of non-labour income. In addition, information on borrowing, inter-family transfers with non-co-resident family members and socio-economic characteristics of non-resident family members was collected for respondent's parents, and up to four siblings and four children. Information on health and education was collected for nearly 3,000 sibling pairs under the age of 15 in wave 1.

Malaysian Family Life Survey (MFLS)	1976/77 (MFLS1) and 1988/89 (MFLS2), 2 waves,	MFLS1 sample consists of 1,262 households with an ever-married woman (aged < 50 years at the time of the initial interview) selected to be representative of Peninsular Malaysia in 1976. 926 of MFLS1 households and a subset of adult children from original households were re-interviewed in MFLS2. In addition, a new sample of 2,184 women age 18–49 (irrespective of their marital status), as well as a sample of 1,357 older Malaysians, age 50 and above were interviewed in 1988/89. Sample sizes of MFLS2 and MFLS3 are 1,239 and 1,207, respectively.	MFLS2 was re-interviewing original MFLS1 respondents and their adult children. MFLS2 consists of four samples: (a) panel sample; (b) children, (c) new members; (d) older respondents. The panel sample consists of original respondents still living in Peninsular Malaysia. The children sample consists of selected children aged 18 and above of original sample members. Interviews were conducted with one child, selected at random, still living in same HH than panel respondents, and a maximum of two children, selected at random, living outside the original HH. The new member sample consists of women aged 18–49 (irrespective of marital status) or ever married women younger than 18 in 1988. The older respondent sample consists of people aged ≥ 50, where one senior per HH was interviewed.	Family structure, fertility, economic status, education, occupations, earnings, transfers, migration, property, gifts, inheritance.	Information on family economic resources such as family income, wealth and economic value of children to parents. Information about types, amounts and direction of transfers during previous 12 months between respondents and relatives.
Mexican Family Life Survey (MxFLS1)	MxFLS1 is a nationally representative survey of individuals, households and communities. The sample consists of over 8,400 households in 150 communities across Mexico. MxFLS2 fieldwork began in 2005 and is scheduled for release in 2007. MxFLS3 will be conducted three years later, in 2008. MxFLS questionnaires follow the design of the IFLS, adapted to the Mexican context.	Retrospective information on formal and informal credits and loans, schooling, employment, mobility and fertility.	Information on income, expenditure, wealth, health, nutrition, education, employment, migration, intra-household allocation and data linking.		

For more detailed information, about the surveys, see Frankenberg and Thomas (2000) for the IFLS, Butz and Da Vanzo (1978) for the MFLS and <http://www.radix.uia.mx/ennvih/main.php?lang=en> for the MxFLS.