Session Number: 2B
Session Title: IntergenerationalTransmission of Inequality
Paper Number: 6
Session Organizers: Miles Corak and Dean Lillard
Discussant: Miles Corak

Paper Prepared for the 28th General Conference of
The International Association for Research in Income and Wealth
Cork, Ireland, August 22 – 28, 2004

Intergenerational Mobility and Sample Selection in Short Panels

Marco Francesconi and Cheti Nicoletti

For additional information please contact:

Cheti Nicoletti
ISER, University of Essex, Wivenhoe Park, CO4 3SQ Colchester, UK
nicolet@essex.ac.uk
Fax 0044 (0) 1206 873151
Tel 0044 (0) 1206 873536

This paper is posted on the following websites: **http://www.iariw.org**
http://www.econ.nyu.edu/iariw
http://www.cso.ie/iariw/iariwhome.html

Intergenerational Mobility and Sample Selection in Short Panels

Marco Francesconi and Cheti Nicoletti
ISER, University of Essex
Wivenhoe Park, CO4 3SQ Colchester, UK
Tel. 0044 1206 873536, E-mail nicolet@essex.ac.uk

July 2004

Abstract

Using data from the first eleven waves of the BHPS, this paper measures the extent of the selection bias induced by adulthood and coresidence conditions — bias that is expected to be severe in short panels — on measures of intergenerational mobility in occupational prestige. We try to limit the impact of other selection biases, such as those induced by labour market restrictions that are typically imposed in intergenerational mobility studies, by using different measures of socioeconomic status that account for missing labour market information. We stress four main results. First, there is evidence of an underestimation of the true intergenerational elasticity, the extent of which ranges between 10 and 25 percent. Second, the proposed methods used to correct for the selection bias seem to be unable to attenuate it, except for the propensity score weighting procedure, which performs well in most circumstances. This result is confirmed both under the assumption of missing-at-random data as well as under the assumption of not-missing-at-random data. Third, the two previous sets of results (direction and extent of the bias, and differential abilities to correct for it) are also robust when we account for measurement error. Fourth, restricting the sample to a period shorter than the eleven waves under analysis leads to a severe sample selection bias. In the cases when the analysis is limited to four waves, this bias may range from 27 to 60 percent.

Key Words: Sample selection; Censored data; Panel data; Intergenerational mobility; Occupational prestige.
JEL classification codes: C23; C24; J24; J62.

1 Introduction

Opinion among economists about the extent of intergenerational mobility has markedly changed in the last twenty years. Twenty years ago, the general view — based primarily on data from the United States — was that there was little persistence in economic status across generations (Becker and Tomes, 1986). A more recent wave of studies in the United States and several other industrialized countries has questioned that view, and found evidence suggestive of far less mobile societies than was earlier believed.[1] Two important problems, which marred early studies, have been underlined. First, the samples used in most analyses were nonrepresentative. In particular, they tended to be of small size and refer to highly homogeneous groups of the population of children and parents (e.g., individuals were from a specific region or city or they were twins). Second, long-run permanent economic status was poorly measured. Most of the early studies used single-year or other short-run measures of economic status, generally proxied by earnings or income.

To limit such problems, the more recent literature uses longitudinal household data from national probability samples. These data tend to avoid sample homogeneity (but may still lead to relatively small estimating samples). Moreover, if they have a long enough time series component, they allow researchers to compute better measures of long-run status, which are typically obtained after averaging data over several years. For example, using five-year averages, Solon (1992) and Zimmerman (1992) estimate intergenerational correlations that are about twice as large as those found in the earlier research surveyed by Becker and Tomes (1986). Using even more years of data, the estimates reported in Mazumder (2001) reveal intergenerational correlations that are about 50 percent greater than those reported in the Solon's and Zimmerman's studies. Recent research has also considerably improved our understanding of the mechanisms that link one generation to the next. Several studies have modeled and estimated the effect on intergenerational mobility of family background (Shea, 2000), neighbourhood influences (Page and Solon, 2003) and assortative mating (Chadwick and Solon, 2002).

However, researchers have virtually ignored one important issue that may plague all intergenerational correlation estimates, namely the issue of selection. The extent of intergenerational mobility is usually measured by estimating a relationship between a measure

[1] See Solon (1999) for a comprehensive review of this more recent literature.

of son's or daughter's economic status (e.g., earnings or income) and the same measure of economic status for his or her parent(s). Common practice has been to exclude all records of data where parents or children report no earnings or income (because, for example, they were unemployed at the time of the survey). Two exceptions are the studies by Couch and Lillard (1998) and Minicozzi (2003). Both conclude that there exists an important role for assumptions on labour market selection in identifying intergenerational income mobility, but their evidence is mixed. Couch and Lillard assign one dollar of income to individuals who have a valid report of no earnings, and find that more selected samples lead to higher correlations between sons' and fathers' incomes (i.e., *less* mobility). Minicozzi uses a different method and estimates different Manski-type bounds around children's income. Contrary to Couch and Lillard, she finds that dropping both unemployed and part-time employed sons leads to a *higher* degree of mobility than if part-time employed sons had been included.

In interegenerational mobility analyses, however, there may be also selection problems that are induced by restrictions not driven by unemployment, part-time employment or lack of labour market information. In the panel data usually used in such analyses, parents and children must be found living in the same household for some time during the panel years and followed over time. Information on children's or parents' status may be censored, for example, by non-ignorable attrition or by the fact that the time series component of longitudinal datasets is not long enough. The aim of this paper is precisely to assess how severe the problem of short panels can be for the estimation of intergenerational mobility between sons' and fathers' socioeconomic status. In short panels, in fact, the choice of the child generation is typically constrained by the trade-off between younger and older children. The former group is likely to be a random sample of children who coreside with their parents but their observed socioeconomic status is almost certainly a noisy measure of long-run status (the condition underlying the choice of this group is what we call *adulthood condition*). The latter group will have better measures of status but its inclusion may bias the sample towards children who coreside with their parents at late ages (this imposes what we call *coresidence condition*).

To address the problem induced by these two conditions we introduce a new methodology which we apply to a short panel using the first eleven waves of the British Household Panel Survey (BHPS), covering the period 1991-2001.[2] All BHPS respondents aged 16 or more are

[2]In this paper we specifically focus on (and model) this double selection problem only, and abstract from

asked to report the occupation of their parents when they were aged 14. This information is based on retrospective questioning (and therefore may suffer from recall error), but it relies neither on the adulthood condition nor on the coresidence condition. Using a continuous index of occupational prestige — which relates strongly to labour income — we then estimate intergenerational elasticities in occupational prestige that are free of selection bias. However, by linking children to parents over the eleven years of the panel and imposing both adulthood and coresidence conditions, we obtain a new selected subsample of son-father pairs. For this subsample, we again estimate intergenerational elasticities. Comparing these elastisticities to those previously estimated provides us with a direct measure of the extent of the selection bias we are interested in. We then evaluate two approaches that correct for this potential bias. The first belongs to the general class of Heckman-type sample selection corrections, which we perform using both parametric and semi-parametric estimation procedures. The second is within the class of models based on propensity score estimation, which we estimate using procedures based on inverse propensity score weighting and on dummies for different levels of the propensity score. Finally we conduct some sensitivity analyses to assess whether the elasticities estimated on the selected sample are robust to varying the length of the panel or the age range of the children.

The rest of the paper is organised as follows. Section 2 describes the issue of selection we are concerned with in this study and casts our contribution within the relevant literature. Section 3 presents the data source, the estimating samples, our measures of socioeconomic status, and the other variables used in estimation. Section 4 discusses our methodology to gauge the extent of selection bias in our sample and to correct for it. Section 5 reports our results and shows a number of sensitivity checks. Section 6 concludes.

2 The Issue

In analysing the extent of intergenerational mobility, researchers have used several different measures of long-run socioeconomic status (e.g., educational attainment, labour income,

the selection issues related to labour market conditions, attrition or nonresponse. However, we will try to account for the issues related to nonresponse and/or out-of-labour-market conditions in ways similar to those used by Couch and Lillard (1998). An empirical analysis of intergenerational mobility where all these different types of selection are jointly modeled and fully accounted for bears investigation in future work.

hourly wages and occupational indices). Each measure has advantages and disadvantages.[3] Because of restrictions imposed by our data, we will use an index of occupational prestige computed according to the technique proposed by Goldthorpe and Hope (1974), which is widely used especially in the sociological literature.[4] There are also several popular methods of examining intergenerational correlations in socioeconomic status. In this paper, we perform our analysis using a log-linear regression model which assumes that the log of a son's socioeconomic status in family i, denoted by y_i, can be expressed as a linear function of his fathers log permanent socio-economic status, x_i, according to

$$y_i = \alpha + \beta x_i + u_i, \tag{1}$$

where β is our parameter of interest that denotes the intergenerational elasticity of son's status with respect to father's status; α is the intercept term that represents the change in status common to the son's generation; and u_i is a random disturbance. Assuming (for simplicity and momentarily) that the underlying variances in father's and son's status are equal (so that β coincides with the correlation between father's and son's status), a value of $\beta = 1$ indicates a situation of complete immobility, whereby (apart from the influence of u_i) the sons' position in their status distribution is fully determined by their fathers' position. A value of $\beta = 0$ instead indicates a situation of complete mobility or regression to the mean, whereby the child's position is completely independent of his father's; and with intermediate values of β between 0 and 1, the status distribution still regresses to the mean, but at a rate that decreases the higher is β.[5]

If y_i and x_i are observed for all sons and fathers in the sample, then ordinary least squares (OLS) estimation will produce a consistent estimate of β. If, however, either y_i or x_i are observed only for some father-son pairs because of selection, the OLS estimate of β obtained from these censored data is inconsistent. There are a number of reasons

[3]For a review, see the discussions in Atkinson, *et al.* (1983), Bowles and Gintis (2002) and Erikson and Goldthorpe (2002).

[4]Goldthorpe and Hope (1974) suggest that the scale which results from their occupational prestige grading exercise should not be viewed as a grading of social status *stricto sensu*, i.e., as tapping some underlying structure of social relations of "deference, acceptance and derogation" (p. 10). It should instead be viewed as "a judgement which is indicative of what might be called the 'general goodness' or ... the 'general desirability' of occupations" (pp. 11-12). Ermisch and Francesconi (2003) present a measurement model that shows the conditions under which the estimates based on the Hope-Goldthorpe index are consistent estimates of parameters similar to those in the economic literature.

[5]It is possible of course to have a negative β (i.e., reversal). In these circumstances, fathers with status above (or below) the mean will have sons with below (or above) average status levels.

why information on x-y pairs may be censored. First, observations where a parent or a child have no required status information are usually deliberately dropped from the analysis. For example, by excluding fathers and sons that were not in full-time jobs (i.e., working on average 30 hours per week at least 30 weeks per year), Zimmerman (1992) implicitly assumes exogenous selection into full-time employment. Solon (1992), Dearden *et al.* (1997), Couch and Dunn (1997) and Chadwick and Solon (2002), among others, exclude cases in which income observations are nonpositive, and thus implicitly assume exogenous selection into employment. These assumptions are not consistent with standard economic results according to which selection into the labour force or into full-time employment is likely to be correlated with potential earnings (Heckman, 1979; Vella, 1998).[6]

Second, information on x-y pairs may be censored because information on child's or father's status is missing (item nonresponse) or because one of the two individuals is not in the sample (unit nonresponse). By dropping father-child observations in which at least one has missing status information or has attrited out of the survey under analysis, most studies implicitly assume that nonresponse is random or that attrition is ignorable.[7] Again, these assumptions run counter to the findings that nonresponse is likely to be systematically correlated with observable and unobservable characteristics, or that estimates of intergenerational relationships are likely to be affected by differential attrition (Fitzgerald *et al.*, 1998).

Third, many longitudinal studies typically restrict their analysis to children from specific birth cohorts, say, born between years b_1 and b_2, $b_1 < b_2$ (e.g., Solon, 1992; Couch and Dunn, 1997; Bjorklund and Jantti, 1997; Chadwick and Solon, 2002; Minicozzi, 2003). The restriction on b_2 (which we label *adulthood condition*) is motivated by the need to assure that children's socioeconomic status is observed as far out as possible on their life cycle (when they are "adult") so that their observed status can be taken as a reliable measure of long-run permanent status.[8] This need has driven researchers to choose a relatively low b_2. The restriction on b_1 (which we label *coresidence condition*) is instead motivated by the need to avoid overrepresenting children who left home at later ages. If there are unobserved

[6] This is the type of selection examined by Couch and Lillard (1998) and Minicozzi (2003).

[7] Selection based on the availability of other relevant variables for children and parents (e.g., age, hours of work, education) will induce similar problems. See Couch and Lillard (1998) for a review of studies that used sample selection criteria on such variables.

[8] Clearly the selection generated by the adulthood condition and the selection induced by the child's labour market conditions described above are closely related. Our empirical analysis will in part address this issue.

factors affecting children's later socioeconomic status (through u_i in (1)) that also influence children's chances of living (or "coresiding") with their parents, then the OLS estimate of β will suffer of sample selection bias. So the need to limit this bias has led researchers to choose a relatively high b_1. This double choice produces the selection problem which this paper focuses on.

In longitudinal birth cohort datasets that follow individuals born in one specific week over time (such as the National Child Development Study used by Dearden *et al.* (1997)), the problem of choosing b_1 and b_2 does not exist since b_1 and b_2 are equal and coincide with the year of the start of the survey, b_0. The drawbacks of such data sources however are that intergenerational analyses can be done only after a relatively long period of time since the surveys started (which may lead to serious attrition bias), and generalisations to other (younger or older) cohorts are not straightforward. In nationally-representative family-based longitudinal data (such as the Panel Study of Income Dynamics used by Solon (1992) or the German Socioeconomic Panel used by Couch and Dunn (1997)), the problem exists and may be serious, especially when $b_1 < b_2 \leq b_0$. For example, in Solon (1992) b_1=1951 so that children are aged at most 17 in 1968 (i.e., b_0, the first year of the survey analysed by Solon) to address the coresidence condition, and b_2=1959 so that children are aged at least 25 in 1984 (the last year of observation) to address the adulthood condition. Apart from attrition issues, this problem disappears if the panel data have a long enough time series component, so that $b_2 \geq b_1 \geq b_0$. But for analyses based on short panels, such as that used in this paper, the problem is potentially severe. Our two main objectives therefore are (a) to quantify the extent to which standard estimates of β obtained from short panels are affected by this bias, and (b) to evaluate alternative econometric techniques to alleviate its effects. Before explaining the methodology to pursue such objectives, in the next section we describe the data used in our empirical application.

3 Data

The data we use are from the first eleven waves of the British Household Panel Survey (BHPS) collected over the period 1991-2001. Since Autumn 1991 the BHPS has annually interviewed a representative sample of about 5,500 households covering more than 10,000 individuals. All adults and children in the first wave are designated as original sample members. On-going representativeness of the non-immigrant population has been maintained by using a "following rule" typical of household panel surveys: at the second and subsequent waves, all original sample members are followed (even if they moved house or if their households split up). Personal interviews are collected, at approximately one-year intervals, for all adult members of all households containing either an original sample member, or an individual born to an original sample member.[9] The sample therefore remains broadly representative of the population of Britain as it changes over time.[10] ¿From the BHPS, we select three different samples and employ various measures of the status variable for sons and fathers, with the double aim of gauging the extent of selection bias of interest here and of attenuating the measurement error problem inherent in all intergenerational studies. We now turn to describe samples and variables.

3.1 Samples

Our main analysis is restricted to 2691 men (or sons) born between 1966 and 1985, who have at least one valid interview over the panel period under study. This represents our Full Sample. The BHPS asks all adult respondents (aged 16 or more) to provide information about their parents' occupations when they (the respondents) were aged 14, and releases data on an index of occupational prestige introduced by Goldthorpe and Hope (1974). This index ranges from 5 to 95, with greater values indicating higher occupational prestige, and it is highly correlated with earnings.[11] Between 1991 and 2001, the BHPS data indicate a correlation between gross monthly earnings and the Hope-Goldthorpe (HG) index of 0.70 for men and 0.75 for women. Because the position of individuals in the occupational hierarchy is relatively stable over time, the HG scale is also likely to be an adequate measure of people's permanent socio-economic status (Nickell, 1982; Ermisch and Francesconi, 2003).

All the individuals in the Full Sample who could be successfully matched to their father are part of our second sample, which we refer to as Restricted Sample. There are 1114 of

[9]Individuals are defined as "adult" (and are therefore interviewed) from their sixtheeth birthday onwards.

[10]Of the individuals interviewed in 1991, 88 percent were re-interviewed in wave 2 (1992). The wave-on-wave response rates from the third wave onwards have been consistently above 95 percent. See Taylor (2003) for a full description of the dataset. Detailed information on the BHPS can also be obtained at <http://www.iser.essex.ac.uk/bhps/doc>. The households from the European Community Household Panel subsample (followed since the seventh wave in 1997), those from the Scotland and Wales booster subsamples (added to the BHPS in the ninth wave) and those from the Northern Ireland booster subsample (which started in wave 11) are excluded from our analysis.

[11]Phelps Brown (1977) reports a strong log-linear relationship between median gross weekly earnings and the HG score, with a rise of 1 unit in the index being associated with an increase of 1.031 percent in earnings. Nickell (1982) finds a correlation between the HG score and the average hourly earnings of 0.85.

such father-son pairs. Differently from the Full Sample, this imposes stringent adulthood and coresidence conditions. Individuals born in 1966 were aged 25 in the first year of the panel (1991) and 36 in the last (2001): they could have lived with their parents at any age between those two years. With a median home-leaving age of about 23-24 (Ermisch and Di Salvo, 1997; Ermisch and Francesconi, 2000), coresidence at such ages means that the Restricted Sample overrepresents sons who left home at late ages. At the other extreme, individuals born in 1985 were aged 16 in 2001 (the last year of analysis): although they are likely to be a random sample of young people living with their parents, their HG index is arguably a noisy measure of long-run status. The comparison of the intergenerational elasticities obtained from the Full Sample and the Restricted Sample will provide us with a measure of the extent of the selection in short panels, under the maintained assumption that Full Sample estimates do not suffer from selection bias.

As discussed in the Introduction, one of the issues emphasised by recent empirical studies is the lack of reliable measures of fathers' long-run permanent status. In this respect our dataset is no exception (see the next subsection). To gauge the impact of the related bias, we thus analyse a third sample (of fathers). In the second (1992) and eleventh (2001) waves of the BHPS, adult respondents were asked to provide information on all their children regardless of where they lived. This, which we call Supplemental Sample, is composed of 1434 men whose sons were born between 1966 and 1985 (the same year-of-birth selection used in the Full Sample). This should provide us with a random sample of fathers with sons born in those years.[12] As measurement error corrections would typically require information on the variance of fathers' permanent status (Zimmerman, 1992), the Supplemental Sample allows us to estimate any moment of distribution of fathers' HG scores, both unconditional and conditional on son's age.

3.2 Socioeconomic Status Variables

We use two alternative measures of sons' socioeconomic status. Assuming exogenous selection into the labour market, the first measure (labelled HG_1^s) is given by the average HG score over all waves after excluding the cases with missing status information either because

[12] As opposed to the Full Sample, however, the Supplemental Sample does not contain information on sons' HG scores, except for those cases in which — as in the Restricted Sample — father and son were observed to live in the same household at least once over the sample period.

the son does not work or because his information is genuinely missing.[13] The second measure (HG_2^s) is given by the average HG score over all waves after replacing the cases in which the son is not in paid employment (except for those who are in full-time education) with the minimum HG score observed in the sample.[14] Table 1 shows that the mean values of HG_1^s and HG_2^s are 43 and 41 respectively for sons in the Full Sample, and 40 and 39 for sons in the Restricted Sample. Regardless of the sample, therefore, HG_2^s is smaller than HG_1^s (as we replaced the cases with missing prestige information with the minimum HG observed), but their differences are not statistically significant. Imputed values, in fact, amount only to 9 percent in the case of the Full Sample and 5 percent in the case of the Restricted Sample.

As mentioned above, one of the major difficulties in estimating intergenerational elasticities abides in the fact that father's status, x_i, is measured with error. The key problem is the lack of direct measures of permanent status. In the case of the Full Sample, in particular, the BHPS provides us with only one single-year measure of fathers' occupational prestige (when sons were aged 14).[15] Although the HG index is an arguably good proxy for long run status, a single-year measure may still be tainted by transitory fluctuations in fathers' careers. In addition, the BHPS elicits this information by asking respondents to report their parents' occupation when they were aged 14.[16] The retrospective questioning of children to obtain data on parents may of course generate recall errors. Both types of errors (due to measurement and recall) may be such that the variance of observed status is greater than the variance of permanent status, leading the OLS estimate of β in (1) to be biased downward.[17] Estimation of (1) is further complicated by the fact that fathers' and sons' occupations may refer to different points in their life cycle. Although this age variation could also bias the estimates of intergenerational mobility, no general conclusion on the direction of this bias can be reached (Jenkins, 1987).

[13] Several studies have argued that averaging status data over time reduces the impact of the transitory component of the status variable (thus reducing the potential of errors-in-variables bias) and yields a more accurate measure of permanent status. See, among others, Solon (1992), Zimmerman (1992), Dearden *et al.* (1997), and Mazumder (2001).

[14] For a similar treatment of non-random selection into the labour force, see Couch and Lillard (1998).

[15] In the case of the Restricted Sample, instead, multiple measures of fathers' HG scores are available in principle. For that sample, however, the measurement problem is aggravated by the fact that fathers and sons are observed at different points in their life cycle (see below). Clearly, the Restricted Sample is expected to suffer from the selection bias discussed in Section 2.

[16] Again, our Restricted Sample does not have to face this problem as fathers and sons report independently their own occupational information.

[17] Because our Supplemental Sample is not affected by such errors, it should give us an idea of how serious this problem might be for the Full Sample.

To address at least part of such problems we use three different measures of fathers' socioeconomic status. The first uses the single-year measure reported by sons, which identifies fathers' occupational prestige when their sons were aged 14, and excludes all cases with missing information (because either the son refused to answer, or he did not know his father's occupation, or his father did not work when he was 14). This is the measure available for the Full Sample, HG_1^f. For all sons in the Full Sample who do not report the occupational information on their fathers but could be successfully matched to them, we replace the missing values on paternal occupational prestige with the HG scores that are available for fathers and average them over all available waves (as long as they are nonmissing). Of course, this replacement can be done only for those father-son pairs that are in the Restricted Sample. We denote this measure HG_2^f. Our last measure, HG_3^f, replaces the missing values of HG_2^f with the minimum score observed in the Full Sample.[18] In the case of the Supplemental Sample we compute two measures of fathers' HG scores. The first, HG_{SS}^f, is the index obtained from the father in the second or eleventh wave when he is asked to report information about all his children. The second, which we denote $\bar{HG}_{SS}^f$ is the average of the HG scores across all waves in which the father is observed. If this information is missing, we replace it with the occupational prestige reported for the most recent job. Both averaging across waves and replacing missing data are meant to reduce the problems induced by missing or inapplicable cases and measurement error (see Section 5).

Table 1 shows that the distributions of HG_1^f and HG_2^f are relatively similar in the Full and Restricted Samples, with means ranging between 47.6 and 49.9 and standard deviations ranging between 14.7 and 15.7. The mean prestige scores in the Supplemental Sample are remarkably close to the scores in the Full Sample, which suggests that measurement and recall errors affecting the status measures in the Full Sample may not be substantial. In addition, we never reject the hypothesis that the mean HG scores and their standard deviations in the Full and Restricted Samples differ for sons or for fathers, even at significance levels as high as 10 percent. The story is different for HG_3^f. For this measure we find lower values in the Restricted Sample at about 46 points, and especially in the Full Sample with a mean score of about 38.6 points. These two values differ at any statistical level, and, within sample, each of them also differs from the other two measures, HG_1^f and HG_2^f. Table 1 shows that the imputations underlying HG_3^f lead to 12- and 54-percent bigger samples in the case of

[18]The rationale for this recoding exercise is the same as that adopted in the case of sons (see above).

the Restricted and Full Samples respectively as compared to the analysis based on HG_2^f. Replacing a large proportion of missing/inapplicable cases (particularly in the Full Sample) may alleviate one problem of selection but may introduce another problem related to how the imputation is performed. This trade-off must be kept in mind while interpreting our results.

3.3 Other Variables

Table 1 lists the summary statistics of the other variables used in the analysis. As in several other studies, the model for son's status in equation (1) is extended to include son's age and its square (Solon, 1992; Zimmerman, 1992; Dearden *et al.*, 1997; Couch and Dunn, 1997; Corak and Heisz, 1999). Although the age range of 16 to 34 years is the same both samples, the mean age of sons in the Full Sample is 23.2, about 2 years greater than in the Restricted Sample. We do not have information on fathers' age in the Full Sample, but in the Restricted Sample their mean age is 48.2, while it is 46.9 in the Supplemental Sample, whose higher standard deviation reflects its greater heterogeneity.

The next section will illustrate our methods to account for sample selection biases. A number of other variables are used for that purpose, and are reported in Table 1. Approximately 94 percent of the sons in both samples are white. The average year of birth for sons in the Full Sample was 1974, compared with 1976 for those in the Restricted Sample. Religious activity is another factor that is believed to have a deep effect on the likelihood of young people's leaving parental home (Cherlin, 1992). Although religious views on family formation are varied, strong religious beliefs are one cultural source of ideas that encourages the maintenance of traditional values (Wilcox, 2002). For each of the three religious denominations considered here (Catholic, Protestant, and other religions), being "active" is defined as attending religious services at least once a month.[19] In both samples, less than 20 percent of the young men are religiously active, more than 40 percent have no religious affiliation and the rest have an affiliation but do not attend services regularly.

Many studies of household formation have underlined the importance of the price of hous-

[19]The "Protestant" group includes: Church of England (Anglican), Church of Scotland, Free Presbyterian, Episcopalian, Methodist, Baptist, Congregationalist, and other Christian denominations. The "other religions" include: Muslim, Hindu, Jewish, Sikh, and other non-Christian denominations. The omitted category includes those with no religious affiliation as well as those who have a religious affiliation but attend religious services only infrequently. Distinguishing between such two groups does not change our results.

ing (e.g., Haurin *et al.*, 1994; Ermisch, 1999). House prices indeed can affect the likelihood of observing fathers and sons living together, and so they may determine the selection into the Restricted Sample.[20] In Table 1, the (log) average prices of housing are relatively similar in the Full and Restricted Samples, with their corresponding levels being about 65,000 and 63,000 respectively. These averages mask large differences across regions and over time. Eight regional dichotomous variables are included in the model for the standard regions of Great Britain.[21] These are the same regions used to assign annual house prices, and so identification of the effects of the price of housing on the probability of not living apart from parents is based on the differences in patterns of house price changes among regions. Table 1 shows that the regional distributions of sons are similar across samples.

4 Methods

The first objective of our analysis is to determine the impact of sample selection assumptions on intergenerational mobility estimates. We begin by estimating an extended (still very parsimonious) specification of equation (1) on the Full and Restricted Samples, namely:

$$y_i = \alpha + \beta x_i + A_i' \gamma + \epsilon_i, \tag{2}$$

[20]The price of housing is an ambiguous concept when there are different housing tenures, non-neutral tax treatment of them, and probable imperfections in financial markets (Ermisch and Di Salvo, 1997; Ermisch, 1999). In Britain in 2002, nearly 90 percent of households are either owner-occupiers (68 percent) or "social tenants" (22 percent). The latter primarily includes households who rent their dwelling from local authorities. Social housing is not allocated by price, but by administrative procedures, which give priority to families with children and the elderly. While only a small proportion of all households rent from private landlords, it is a relatively important sector for young people leaving their parental home, being the destination of all 45 percent of all departures and 33 percent of departures among those who are not full-time students. Owner-occupation is the destination for 56 percent of non-student departures. Information on rents in the private market is not available in the BHPS, and there are barriers to entry into social rental housing for young people. Our measure of the price of housing is the same as that used by Ermisch (1999), and is given by the average "mix-adjusted" house price relative to the retail price index in any given year for the region in which a person resided in that year. It adjusts for changes in the mix of the size and type of house (e.g., detached, semi-detached, flat, etc.), but does not adjust for quality change. This measure is likely to capture a large proportion of the variation in a measure of the annual "user cost of housing" for owner-occupiers, because mortgage and income tax rates are set nationally and relative house prices show much larger variation over time than these. It also could be viewed as an indicator of housing market conditions, in both rental and owner-occupied markets. For individuals in the Full Sample who could not be matched with their fathers, the price of housing refers to the price observed in the first wave they were in the panel. For those who coreside with their fathers (and therefore, all sons in the Restricted Sample), this variable is measured at the last wave they were observed living together.

[21]The regions are Greater London (which is our base category), South East, South West, East Anglia and East and West Midlands, North West (including Yorkshire and Humberside), Rest of the North, Wales, and Scotland.

where A_i is a vector containing the son's age and age squared, and ϵ is the new error term. Comparing then the estimates of β obtained from the two samples provides us with a measure of the extent of the bias. We will perform this comparison by using each of the two HG score measures for sons alternatively with each of the three measures for fathers.

After having established the magnitude of the bias (if any), we consider different estimation methods to correct for it. For this purpose, we apply five different methods.[22] The first method is based on a maximum likelihood (ML) estimation of the main equation (2) jointly with the following selection equation:

$$r_i^* = Z_i'\Theta + v_i, \tag{3}$$

with

$$\begin{aligned} r_i &= 1 \text{ if } \quad r_i^* > 0 \\ r_i &= 0 \text{ otherwise,} \end{aligned} \tag{4}$$

where r_i^* is a latent variable with associated indicator function r_i that takes value 1 if the father's status, x_i, is observed, and 0 if x_i is missing; Z_i is a vector of explanatory variables (possibly including A_i) that determine the probability of observing a son living with his father during the survey years and are assumed for simplicity to be observed for all individuals; v_i is a zero-mean error term with $E(\epsilon_i \,|\, v_i) \neq 0$; and Θ is a conformable vector of parameters to be estimated. Under the assumption that ϵ_i and v_i are independently and identically distributed $N(0, \Sigma)$, with Σ being a full-rank variance-covariance matrix, and (ϵ_i, v_i) independent of Z_i, it is straighforward to estimate all the parameters in (2) and (3) using a standard ML procedure.

The second method follows the parametric two-step (TS) estimation procedure suggested by Heckman (1979), which, like the previous method, imposes a bivariate normal distribution on the error terms, ϵ_i and v_i. The idea here is that in the conditional expectation of (2), after normalising σ_v to 1, the term

$$E(\epsilon_i \,|\, Z_i, r_i = 1) = \sigma_{\epsilon v}\Big(\frac{\phi(Z_i'\Theta)}{\Phi(Z_i'\Theta)}\Big),$$

where $\phi(\cdot)$ and $\Phi(\cdot)$ denote the prabability density and cumulative distribution functions of the standard normal distribution, and the term in parentheses is known as the inverse

[22] For a comprehensive review, see Vella (1998).

Mill's ratio. Therefore $E(\epsilon_i \,|\, Z_i, r_i = 1)$ is different from zero as long as the error terms in (2) and (3) are correlated, that is, $\sigma_{\epsilon v} \neq 0$. The TS method is first to obtain an estimate of the inverse Mill's ratio using a probit model , and then obtain a consistent estimate of β (and the other parameters in (2)) by OLS. Without imposing any exclusion restriction, both ML and TS methods will identify our model (2)-(4) via functional form (i.e., through the nonlinearity implied by the error structure). However, in the analysis below we introduce exclusion restrictions, with our selection equation including the same variables in (2) plus dummy variables for region of residence and race and the regional house price index.[23]

A criticism of the ML and TS methods is their reliance on distributional assumptions on ϵ_i and v_i (Lee, 1984; Bera *et al.*, 1984). In particular, bivariate normality implies that the relationship between the error terms is linear. Our third method tries to test for departures from normality by including terms that capture systematic deviations from linearity in the disturbances. This method again involves a two-step procedure, in which the first step is identical to that of TS. The second step estimates the conditional expectation of (2), but — rather than including the first-step estimate of the inverse Mill's ratio as an additional regressor — it includes a polynomial expression that is proportional to the inverse Mill's ratio and takes the form

$$E(\epsilon_i \,|\, Z_i, r_i) = \left[\frac{\phi(Z_i'\Theta)}{\Phi(Z_i'\Theta)}\right] \sum_{j=0}^{J} (Z_i'\Theta)^j,$$

where J is the order of the polynomial. If $J = 0$, this expression equals the inverse Mill's ratio of the TS method. But coefficients of the terms for $j = 1, ..., J$ that are statistically different from zero lead to reject the hypothesis that the disturbances are bivariate normal. In our application J=4. We refer to this method as TSN, to underline its two-step nature and the fact that it tests for nonnormality.

Our last two methods rely less on distributional assumptions by imposing an index restriction according to which $E(\epsilon_i \,|\, Z_i, r_i) = g(Z_i'\Theta) = h(Pr(r_i = 1 \,|\, Z_i))$, where both g and h are unknown functions, and $Pr(r_i = 1 \,|\, Z_i) = \Phi(Z_i'\Theta)$ is the propensity score. The fourth method follows a two-step propensity score stratification (PSS) procedure by slightly modifying the estimator proposed by Cosslett (1991). In the first step we estimate the selection

[23] Similar restrictions are used in Vella (1998). Clearly, for our exclusion restrictions to be valid, we need region, race and the house price index: (a) to be predictive of r_i; (b) to assign the realisation of r_i randomly across families given the other observables; and (c) to be exogenous to ϵ_i.

model parametrically via probit (rather than nonparametrically as suggested by Cosslett). In the second step, we estimate our main equation (2) while approximating the selection correction, $g(\cdot)$, by J indicator variables, $d_j = \{\mathrm{I}(Pr(r_i = 1 \mid Z_i) \in j)\}$, where $j = 1, \dot{,} J$ is one of the J intervals that partition the [0,1] support of the propensity score.[24] In our analysis, we use two alternative specifications of the PSS method. In one, we consider four dummy variables indicating whether the predicted propensity score is in the bottom quartile, or between the 25th percentile and the median, or between the median and the 75th percentile, or in the top quartile. In the other, we consider ten dummy variables that partition the predicted propensity score distribution in deciles.

Finally, our fifth method applies the propensity score weighting (PSW) procedure recently used in Robins and Rotnitzky (1995), Robins *et al.* (1995), and Abowd *et al.* (2001). This is again a two-step method with the first step requiring the usual estimation of the selection model via probit. In the second step, we estimate the main equation (2) using weighted least squares with weights given by the inverse propensity scores obtained from the first step.

A key issue in selection problems has to do with the way in which the selection process itself is believed to be generated. If the probability that $r_i = 1$ is independent on $\{y_i, x_i, Z_i\}$, then the data are *missing completely at random*, the missing data problem can be ignored, and we can estimate equation (2) using the subsample of father-child pairs with observations on both y_i and x_i (and other regressors). If instead r_i is not independent of y_i, but, conditional on y_i and Z_i, it is uncorrelated to x_i, the data are said to be *missing at random* (MAR), and consistent estimates of β can be obtained from (2) using the subsample of individuals with complete observations and considering the missing covariate problem as if it were a missing dependent variable problem.

If r_i is independent neither of y_i nor of x_i, the data are *not missing at random* (NMAR), and consistent estimates of β can be obtained only by inverting (2) into

$$x_i = a + by_i + A_i'c + \xi_i, \tag{5}$$

which allows us to consider the problem of the missing x_i as a standard missing dependent variable problem. In the reverse-OLS regressions (5), a, b, and c are paramaters, and ξ_i is an error term.[25] Since our parameter of interest is β in equation (2), this can be recovered

[24] Although similar to ours, the procedure by Cosslett (1991) suggests a partition of the support of $Z\Theta$ instead of the support $[0, 1]$ of the propensity score.

[25] It is worthwhile noticing that we estimated one specification of the selection model (3)-(4) in relation to

from (5) through[26]

$$\widehat{\beta} = \widehat{b}\Big(\frac{\sigma^2_{\tilde{y},\tilde{y}}}{\sigma^2_{\tilde{x},\tilde{x}}}\Big), \tag{6}$$

where $\sigma^2_{\tilde{k},\tilde{k}}$ is the variance of $\tilde{k}$, $\tilde{w} = \tilde{y}, \tilde{x}$, and $\tilde{w}$ is the residual of the regression of w on A. Now, y_i and A_i are observed for all father-son pairs, so the estimation of $\sigma_{\tilde{y},\tilde{y}}$ is unproblematic. However, x_i is observed only for father-child pairs for whom both the adulthood and the coresidence conditions hold, the possibility of missing x_i's must be accounted for in estimating $\sigma_{\tilde{x},\tilde{x}}$. For this purpose, the information contained in the Supplemental Sample is crucial. Specifically, for all fathers in that sample, we regress the log of their HG index on a constant, their son's age and age squared, and use the residuals of this regression to compute $\sigma_{\tilde{x},\tilde{x}}$. Therefore, for each of the five methods described above (ML, TS, TSN, PSS, and PSW), we will present two sets of estimates of β, one set that assumes data missing at random and another set that assumes data not missing at random.

Table 2 summarises the assumptions imposed by each of the methods described in this section to obtain a consistent estimate of the intergenerational correlation parameter.[27] For example, assumption A1, which is needed for the OLS estimator of β in (2) to be consistent, requires y_i to be independent of r_i conditional on x_i and A_i. This means that the residuals in (2) must not depend on the selection process. However, such a dependency could be driven either by unobserved variables — in which case the ML, TS, TSN, and PSS methods and their corresponding assumptions will apply — or only by observed variables that are excluded from the main equation (2). In this latter case, assumption A5 becomes relevant.[28] Finally, assumption A3 identifies our notion of MAR, which naturally will be relaxed when we estimate our models under the hypothesis of NMAR data.

(2) and another specification in relation to (5). In the first, the Z_i vector excludes both y_i (obviously because it is the dependent variable) and x_i (because we impose MAR). In the second specification, we exclude x_i (since it is the dependent variable in (5)) but we include y_i.

[26]For notational convenience, we drop the subscript i.

[27]The notation "$\perp\!\!\!\perp$" indicates statistical independence.

[28]A similar reasoning is valid when the reverse-regression model (5 is used.

5 Results

5.1 The Extent of Selection Bias

Table 3 reports the estimated intergenerational elasticities using six different combinations of the son's and father's status variables for the Full and Restricted Samples computed under the assumption that y_i is independent of r_i conditional on all the explanatory variables in (2). The estimates from the Full Sample range between 0.08 and 0.23, while those from the Restricted Sample range between 0.11 and 0.20.[29] The last column of Table 3 reveals that the gap between the two sets of estimates is statistically significant at conventional levels, except for the case in which y_i is proxied by HG_2^s and x_i by HG_1^f (second row).[30]

Taking the values from the Full Sample as benchmark estimates free of selection bias, we notice that the direction of the bias depends on whether we use HG_3^f as a measure of father's status. As discussed in subsection 3.2, this allows for larger sample sizes but at the cost of replacing more than 50 percent of the missing observations on x_i with the minimum observed HG score. When HG_3^f is used, the elasticity estimates from the Restricted Sample are always greater than those from the Full Sample. This is consistent with the findings by Couch and Lillard (1998). The difference is small but still significant in one case (sixth row of Table 3), and very large, of the order of 46 percent, in the other case (fifth row). However, this overestimation is likely to be the result of the imputation procedure used to construct HG_3^f. In the other four cases (i.e., when HG_1^s, HG_2^s, HG_1^f, and HG_2^f are used), instead, we always find that the Restricted Sample leads to an underestimation of the true intergenerational elasticity, which is in line with the results obtained in a different context by Minicozzi (2003). The magnitude of the bias varies from moderate (between 12 and 14

[29]In spite of being smaller than the elasticities reported in many recent studies that use earnings or income as measures of status (Solon, 1999), our estimates are close to those shown in Atkinson *et al.* (1983) for Britain, when they use *net* family incomes rather than earnings as their variables of interest. They are also close to those reported in Blanden *et al.* (2004), where the log of children's earnings are regressed on the log of parental income, and to those reported in Ermisch and Francesconi (2004), who also use HG scores.

[30]Using a probit model in which r_i is determined by y_i, x_i and A_i, we can also check if assumption A1 holds by testing whether the coefficient on y_i is significantly different from zero. Differently from the Chow test shown in Table 3, this test does not require linearity, but imposes a parametric probability model for the selection process. The results are reported in Appendix Table 10 (panel (a)). In three out of six cases assumption A1 is rejected (rows one, three, and five). This indicates that selection issues are relevant. Importantly, in only one case (row four), this test contradicts the results from the Chow test. Notice also that, in this case, the normality assumption imposed by the probit specification is always rejected except when HG_1^s and HG_2^f are used (see panel (h)). If, instead, the probit selection model includes Z among its explanatory variables, normality is never rejected.

percent in the first two rows) to large (about 25 percent in the third and fourth rows). Both the direction of the bias and its extent therefore depend on how socioeconomic status is measured.

Biases of the order of 25 percent are arguably large. But even those found in the first two rows of Table 3 are sizeable, in the sense that they are likely to be consequential to mobility. This is because of our measure of socioeconomic status. In general, values of β of 0.18 or 0.19 suggest patterns of intergenerational mobility that are relatively similar to those implied by values of about 0.22 or 0.23. The sons' (conditional) probability of staying in or moving to different points of the HG scores distribution varies by 1 or 2 percentage points as long as their father's prestige does not lie at the extremes of his distribution. However, if father's prestige is at the extremes (e.g., bottom and top deciles), the differences in probability are larger (of the order of 3-4 percentage points) and such differences may underpin substantially different occupations and earnings. For example, almost 90 percent of fathers in the top decile are managers (across all industrial sectors) and professionals (e.g., engineers, architects, university professors, medical doctors, solicitors and chartered accountants). This is true only for 58 percent of fathers in the ninth decile of the HG score distribution (which lies between the 80th and the 90th percentiles). Those occupational differences are reflected in substantial pay differentials. For example, the 2001 average monthly earnings is 2,100 for fathers in the ninth decile, while fathers in the top decile earn 2,500 per month, approximately 15 percent more. In sum, a downward bias of up to 25 percent in intergenerational elasticities based on occupational prestige measures is likely to provide us with a different picture of social mobility even if the "true" value of the elasticity is low.

5.2 Correcting for Sample Selection Bias

The results for the (probit) selection model are reported in Appendix Table 9. We show three specifications. Specification [1] is relevant for equation (2), specifications [2] and [3] pertain to model (5) but use two different measures of y_i, HG_1^s and HG_2^s respectively. Our estimates confirm a number of previous results (e.g., Haurin *et al.*, 1994; Ermisch and Di Salvo, 1997). Individuals who live in Greater London are less likely to coreside with their fathers in any of the panel years as compared to individuals in other parts of the country. Older sons are less likely to be matched with their fathers, and so are young Black (Carribean or African) men. Men of Pakistani or Bangladeshi ethnic origin are instead more likely to be observed living with their fathers than otherwise identical White men. This is also the case for young men who are religiously active (but only if they are non-Christians). Higher house prices significantly reduce the probability of observing father-son pairs, suggesting a higher rate of mismatch in the panel in times or areas with relatively higher house prices. This is in contrast with the findings reported in other studies (e.g., Ermisch, 1999), although such studies were typically interested in the household formation process rather than the coresidence of fathers and sons (indeed a number of sons who could not be matched with their fathers in our sample are found to live with their mothers).

5.2.1 Correcting under MAR Data

We now turn to see how well the methods described in Section 4 reduce the sample selection bias that affects β obtained from the Restricted Sample. Table 4 shows these results under the assumption of MAR data.[31] For expositional convenience, the top panel of Table 4 reports again the elasticities shown in Table 3 for all the possible combinations of HG^s and HG^f. All correction methods seem to be unable to attenuate the bias, except for the PSW method. Excluding the cases in which HG_3^f is used (when the PSW procedure aggravates the upward bias in β induced in the Restricted Sample), in all other cases this method performs very well. For example, in the third and fourth columns (where the bias is 25 and 24 percent respectively), the PSW-corrected values of β are 0.16 and 0.17 respectively, or 12 and 6 percent smaller than the corresponding true values in the Full Sample.

The PSW method is valid when assumptions A3 and A5 are satisfied (see Table 2). The tests reported in Appendix Table 10 panel (c) reveal that, for all models that exclude HG_3^f, assumption A3 cannot be rejected at conventional levels of statistical significance. Assumption A5 instead appears to hold in two cases only (rows two and four in panel (e)). Unsurprisingly, the fact that the correction methods other than PSW perform poorly is confirmed by the rejection of A4 regardless of the measures of socioeconomic status (panel (d)). Taken together, therefore, these results indicate that the selection of interest here is primarily driven by observables.

[31] We conducted a number of tests on the significance of the additional variables needed to correct for sample selection in the ML, TS, TSN and PSS procedures. In all cases, we cannot reject the hypothesis that such variables are not significantly different from zero. However, as the Chow tests reported in Table 3 reveal, these tests seem to be not powerful enough to detect significant differences in intergenerational elasticities.

5.2.2 Correcting under NMAR Data

One procedure to obtain consistent estimates of β when the data are NMAR is: (i) to estimate the revese regression (5); and (ii) to multiply the b coefficient by the ratio in variances as shown in (6). The first two columns of Table 5 report the estimates of b for the Full and Restricted Samples respectively. Typically, these estimates (in either sample) are greater than those found for β in Table 3. This is especially evident in the case of the regressions that use HG_3^f as the measure of father's status. Given the link between b and β, such a result implies that the variability in status among fathers is greater than that among sons (see also Table 1). This relationship is plausible if the variance in status increases over the life cycle (Grawe, 2004). Comparing the estimates in the two samples, we find that the Restricted Sample produces a greater b in two cases out of six (second and fifth rows). However, as revealed in the Chow tests reported in the third column of the table, the differences are significant only for the estimates in the last two rows.[32]

The next two columns of Table 5 show the estimates of the variance in father's status, $\sigma^2_{\tilde{x},\tilde{x}}$, computed on the Full and Supplemental Samples, respectively.[33] We cannot reject the hypothesis that $\sigma^2_{\tilde{x},\tilde{x}}$ under the Full Sample is equal to $\sigma^2_{\tilde{x},\tilde{x}}$ under the Supplemental Sample for the first measure of father's status, HG_1^f. In the other cases, however, the two variances are significantly different at standard levels.

The intergenerational elasticities implied by the estimates of b and $\sigma_{\tilde{x},\tilde{x}}$ according to expression (6) are in Table 6. The first row shows the values of β found when b and both $\sigma^2_{\tilde{x},\tilde{x}}$ and $\sigma^2_{\tilde{y},\tilde{y}}$ are from the Full Sample. This produces the same estimates of β reported in Table 3, which — in what follows — are again taken as our new benchmark. The values in the second and third rows differ by the estimates of b (which are computed on the Full and Restricted Samples, respectively), but use the same ratio of variances, namely the ratio between the variance in son's status computed on the Full Sample and the variance in father's status computed on the Supplemental Sample.

Using HG_3^f yields large overestimates of β (see the last two columns of Table 6). As

[32]With the exception of the case when $HG_2^s and \mathrm{HG}_1^f$ are used, the test for the validity of assumption A2 leads to the same result (see Appendix Table 10, panel (b)).

[33]Notice that the variances for the Full Sample are identical each time we use the same measure of father's status, while those for the Supplemental Sample are invariant to the measure used. We do not report the estimates of the variance for the Restricted Sample, because they will yield the same estimates of β as those shown in Table 4. We also do not report the estimates of the variance in son's status, $\sigma^2_{\tilde{y},\tilde{y}}, since \mathrm{y}_i$ and the other relevant variables in (5) are observed for all father-son pairs.

discussed earlier, this is likely to be the result of our imputation of missing observations on x_i. For the other measures, the Restricted Sample underestimates the true value of β, with the bias ranging from moderate (about 10 percent in the first column) to large (about 19 and 34 percent in the fourth and third columns, respectively), except for the case when HG_2^s and HG_1^f are used. This last case represents the only striking departure from the values reported in Table 3, and implies an upward bias of about 11 percent. Thus, β is biased and in general likely to be underestimated, at least if one disregards our more problematic measure of father's status, HG_3^f.

Table 6 shows that the PSW method corrects for the bias in the first, third, and fourth columns quite well. Interestingly, these are the only models for which assumption A7 is not rejected. The other methods perform less satisfactorily, perhaps because their underlying assumption A6 is always rejected. These findings broadly confirm those obtained above under the hypothesis of MAR data, and again stress the importance of selection on observables.

5.2.3 Measurement Error

As mentioned in the Introduction, a number of studies have emphasised the role played by measurement error in socioeconomic status in estimating intergenerational mobility. The unavailability of measures of permanent status generally leads to downward-biased elasticities (that is, greater mobility). Assuming a classical measurement error model,[34] one way to attenuate the bias has been to average over repeated observations on father's status (Solon, 1992; Zimmerman, 1992). It is straighfoward to show that

$$\beta_{avg} = \beta\left(\frac{\sigma^2_{\tilde{x},tildex}}{\sigma^2_{\tilde{x}_{avg},\tilde{x}_{avg}}}\right), \tag{7}$$

where β is the elasticity obtained from model (2), $\sigma^2_{\tilde{x},tildex}$ is the variance of the residual of the regression of x_{it} on A_{it} for all years t in which fathers are observed, and $\sigma^2_{\tilde{x}_{avg},\tilde{x}_{avg}}$ is the variance of the residual of the regression of time averages of x on A. The problem is that in either the Full or Restricted Samples we do not have repeated observations on father's status. To compute the term at the denominator in parentheses in expression (7), therefore, we have to resort to the "external" information contained in the Supplemental Sample. In

[34]This assumes that the measurement errors for sons and fathers are mutually uncorrelated and also uncorrelated with the permanent component of status (see Fuller, 1987).

particular, our measure of interest is the average of the father's HG scores across all waves in which he is observed, $\bar{HG}^f_{SS}$ (see subsection 3.2).

The results that adjust β to account for measurement error according to (7) are reported in Table 7. The last two columns show elasticities that are 2.5 times higher than those in Table 3 (Full Sample). This difference underlines that the HG^f_3 measure is likely to be contaminated by substantial measurement error. On the other hand, the first two columns show only marginally greater estimates and the third and fourth columns only marginally smaller estimates than those reported in Table 3, suggesting that the HG score measures used here are likely to be good proxies of permanent economic status. Apart from such differences, the findings emerged earlier are still valid. We highlight two of such findings. First, using the Restricted Sample leads to underestimate the true intergenerational elasticity (with biases ranging between 12 and 25 percent), except when HG^f_3 is used (in which cases we obtain upward-biased estimates). Second, all correction methods perform poorly, in the sense that they are unable to attenuate the selection bias, apart from the PSW procedure for the cases in which we detect downward biases. In some instances, the PSW-corrected estimates reduce the bias by an order of 4 (from 25 percent to 6 percent, see the fourth column in Table 7).

5.3 The Effect of Changing the Length of the Panel

In this section we present a sensitivity analysis of the intergenerational mobility when restricting the sample to a period shorter than 11 waves. In Table 8 we report the β estimates computed using three new samples, namely, the subsamples of sons coeresident with their fathers in at least one wave in the first 8, 6 and 4 waves. For each subsample we compute the measures of occupational prestige using only information available in the fictitiously shorter panel period (that is, the first eight, six, and four waves, respectevely). Such measures are not directly comparable to those of the Full Sample, since the latter are computed over 11 waves of data.[35]

Two comments are in order. First, limiting the analysis to a smaller number of waves leads to smaller sample sizes, and this may strongly decrease the estimation precision. While in the restricted sample based on 8 and 11 waves all β are significantly different from zero at both 1 and 5 percent significance levels, in the subsamples based on six and four waves here are a few cases in which we cannot reject the assumption of a zero intergenerational elasticity.

[35]This means that the correction procedures cannot be straighfowardly carried out.

Second, excluding the cases in which HG_3^f is used, the intergenerational elasticity declines as the length of observation shrinks. This has substantial effects on the underestimation of β. In fact, while the downward bias ranges from 12 to 25 percent when we use all eleven waves, it ranges from 27 to almost 60 percent when we use only four waves.

6 Conclusion

Using data from the first eleven waves of the BHPS, this paper measures the extent of the selection bias induced by adulthood and coresidence conditions — bias that is expected to be severe in short panels — on measures of intergenerational mobility in occupational prestige. We try to limit the impact of other selection biases, such as those induced by labour market restrictions that are typically imposed in intergenerational mobility studies, by using different measures of socioeconomic status that account for missing labour market information.

We stress four main results. First, there is evidence of an underestimation of the true intergenerational elasticity, although some more noisy measures of (father's) status provide support for upward-biased estimates. The extent of the downward bias is moderate in some cases (of the order of 10-12 percent) and large in others (of the order of 25 percent). The consequences in terms of intergenerational mobility of such biases are noticeable especially at the extremes of the occupational prestige distribution. Second, the proposed methods used to correct for the selection bias seem to be unable to attenuate it, except for the propensity score weighting procedure, which performs well in most circumstances. In some cases, the PSW-corrected estimates of the intergenerational elasticity reduce the bias by an order of 4 (from 25 percent to 6 percent). This result is confirmed both under the assumption of missing-at-random data as well as under the assumption of not-missing-at-random data. Third, the two previous sets of results (direction and extent of the bias, and differential abilities to correct for it) are also robust when we account for measurement error. Fourth, restricting the sample to a period shorter than the eleven waves under analysis leads to a severe sample selection bias. In the cases when the analysis is limited to four waves, this bias may range from 27 to 60 percent.

References

Abowd J., Crépon B., Kramarz F. (2001) "Moment Estimation with Attrition." *Journal of the American Statistical Association*, 96, 456, 1223-1231.

Atkinson, A.F., Maynard, A.K., Trinder, C.G. (1983) *Parents and Children: Incomes in Two Generations.* London: Heinemann.

Becker, G.S., Tomes, N. (1986) "Human Capital and the Rise and Fall of Families." *Journal of Labor Economics*, 4, S1–S39.

Bera, A., Jarque, C., Lee L.-F. (1984) "Testing the Normality Assumptions in Limited Dependent Variable Models." *International Economic Review*, 25, 563–78.

Blanden, J., Goodman, A., Gregg, P., Machin, S. (2004) "Changes in Intergenerational Mobility in Britain." In M. Corak (ed.) *Generational Income Mobility in North America and Europe.* Cambridge: Cambridge University Press (forthcoming).

Bowles, S., Gintis, H. (2002) "The Inheritance of Inequality." *Journal of Economic Perspectives*, 16, 3–30.

Chadwick, L., Solon, G. (2002) "Intergenerational Income Mobility Among Daughters." *American Economic Review*, 92, 335–44.

Cherlin, A.J. (1992) *Marriage, Divorce and Remarriage.* Cambridge, MA: Harvard University Press.

Corak, M., Heisz, A. (1998) "The Intergenerational Earnings and Income Mobility of Canadian Men: Evidence from Longitudinal Income Tax Data." *Journal of Human Resources*, 34, 504–33.

Cosslett, S. (1991) "Semiparametric Estimation of a Regression Model with Sample Selectivity." In W.A. Barnett, J. Powell, and G. Tauchem (eds.) *Nonparametric and Semiparametric Methods in Econometrics and Statistics*, 175–97. Cambridge: Cambridge University Press.

Couch, K.A., Dunn, T.A. (1997) "Intergenerational Correlations in Labor Market Status: A Comparison of the United States and Germany." *Journal of Human Resources*, 32, 210–32.

Couch, K.A., Lillard, D.R. (1998) "Sample Selection Rules and the Intergenerational Correlation of Earnings." *Labour Economics*, 5, 313–29.

Dearden, L., Machin, S., Reed, H. (1997) "Intergenerational Mobility in Britain." *Economic Journal*, 107, 47-66.

Erikson, R., Godthorpe, J.H. (2002) "Intergenerational Inequality: A Sociological Perspective." *Journal of Economic Perspectives*, 16, 31–44.

Ermisch, J., Di Salvo, P. (1997) "The Economic Determinants of Young People's Household Formation." *Economica*, 64, 627–44.

Ermisch, J. (1999) "Prices, Parents, and Young People's Household Formation." *Journal of Urban Economics*, 45, 47–71.

Ermisch, J., Francesconi, M. (2000) "Patterns of Household and Family Formation." In R. Berthoud and J. Gershuny (eds.) *Seven Years in the Lives of British Families.* Bristol: Policy Press.

Ermisch, J., Francesconi, M. (2003) "Intergenerational Social Mobility and Assortative Mating in Britain." Unpublished manuscript. Colchester: University of Essex, October.

Ermisch, J., Francesconi, M. (2004) "Intergenerational Mobility in Britain: New Evidence from the BHPS." In M. Corak (ed.) *Generational Income Mobility in North America and Europe.* Cambridge: Cambridge University Press (forthcoming).

Fitzgerald, J., Gottschalk, P. Moffitt, R. (1998) "An Analysis of the Impact of Sample Attrition on the Second Generation of Respondents in the Michigan Panel Study of Income Dynamics." *Journal of Human Resources*, 33, 300-44.

Fuller W.A. (1987) *Measurement Error Models.* New York: Wiley.

Grawe, N.D. (2004) "Life-Cycle Bias in the Estimation of Intergenerational Income Persistence." In M. Corak (ed.) *Generational Mobility ...* (forthcoming).

Haurin, D.R., Hendershott, P.H., Kim, D. (1994) "Housing Decisions of American Youth." *Journal fo Urban Economics*, 35, 28–44.

Heckman J. (1979) "Sample Selection as a Specification Error." *Econometrica*, 47, 153–161.

Jenkins, S.P. (1987) "Snapshots versus Movies: 'Lifecycle Bias' and the Estimation of Intergenerational Earnings Inheritance." *European Economic Review*, 31, 1149–58.

Lee L.–F. (1984) "Tests for the Bivariate Normal Distribution in Econometric Models with Selectivity." *Econometrica* , 52, 843–63.

Mazumder, B. (2001) "Earnings Mobility in the US: A New Look at Intergenerational Mobility." Center for Labor Research, University of California, Berkeley, Working Paper No. 34, March.

Minicozzi, A.L. (2003) "Estimation of Sons' Intergenerational Earnings Mobility in the Presence of Censoring." *Journal of Applied Econometrics*, 18, 291–314.

Page, M.E., Solon, G. (2003) "Correlations between Sisters and Neighbouring Girls in Their Subsequent Income as Adults." *Journal of Applied Econometrics*, 18, 545–62.

Robins J., Rotnitzky A. (1995) "Semiparametric Efficiency in Multivariate Regression Models with Missing Data." *Journal of the American Statistical Association*, 90, 122–129.

Robins J., Rotnitzky A., Zhao L. (1995) "Analysis of Semiparametric Regression Models for Repeated Outcomes in Presence of Missing Data," *Journal of the American Statistical Association*, 90, 106-121.

Shea, J. (2000) "Does Family Income Matter?" *Journal of Public Economics*, ..

Solon, G. (1992) "Intergenerational Income Mobility in the United States." *American Economic Review*, 82, 393–408.

Solon, G. (1999) "Intergenerational Mobility in the Labor Market." In O. Ashenfelter and D. Card (eds.) *Handbook of Labor Economics*, Volume 3, Chapter 29, 1761–1800. Amsterdam: Elsevier.

Taylor M.F. (ed.) with Brice, J., Buck, N., Prentice-Lane, E. (2002) *British Household Panel Survey User Manual Volume A: Introduction, Technical Report and Appendices.* Colchester: University of Essex.

Vella F. (1998) "Estimating Models with Sample Selection Bias: A Survey." *Journal of Human Resources*, 3, 127–169.

Wilcox, W.B. (2002) "Then Comes Marriage? Religion, Race, and Marriage in Urban America." Working Paper 2002-13-FF. Princeton, NJ: Center for Research on Child Wellbeing, Princeton University.

Zimmerman, D.J. (1992) "Regression Toward Mediocrity in Economic Stature." *American Economic Review*, 82, 409–29.

Table 1: Summary Statistics (mean out of and standard deviation in parenthesis).

Variables	Full sample		Restricted sample		Fathers' sample	
Variables	Statistics	No. obs.	Statistics	No. obs.	Statistics	No. obs.
Hope-Golthorpe score for sons						
HG_1^s	42.99 (13.33)	2117	40.00(11.63)	969		
HG_2^s	40.92 (14.47)	2311	38.88 (12.30)	1022		
Hope-Golthorpe score for fathers						
HG_1^f	49.63 (15.70)	1152	47.57 (15.14)	397		
HG_2^f	49.91 (15.24)	1753	49.29 (14.72)	998		
HG_3^f	38.62 (19.74)	2691	45.99 (16.98)	1114		
HG_{SS}^f					47.80 (15.86)	1427
$\bar{HG}_{SS}^f$ mean					47.68 (15.01)	1430
Other characteristics						
Father's age			48.21 (7.54)	1062	46.93 (8.17)	1434
House price index	11.08 (0.26)	2691	11.05 (0.32)	1114		
Sons' characteristics						
Age	23.15 (4.51)	2691	21.30	1114	19.76 (6.64)	1434
Year of birth	1974 (5.32)	2691	1976 (5.20)	1114	1976 (5.80)	1434
White	0.94	2691	0.94	1114		
Black	0.02	2691	0.01	1114		
Indian	0.02	2691	0.02	1114		
Pakistani/Bang.	0.01	2691	0.02	1114		
Other race	0.01	2691	0.01	1114		
Active Catholic	0.06	2691	0.05	1114		
Active Protestant	0.12	2691	0.11	1114		
Active other religion	0.02	2691	0.03	1114		
Region of sons' residence						
London	10.52%	2691	9.96%	1114		
South West	8.62%	2691	8.44%	1114		
Rest of South East	18.62%	2691	19.30%	1114		
Anglia and Midlands	23.67%	2691	23.42%	1114		
North West	10.14%	2691	11.13%	1114		
Rest of North	16.31%	2691	16.97%	1114		
Wales	4.42%	2691	4.76%	1114		
Scotland	7.69%	2691	6.01%	1114		

Table 2: Assumptions imposed by different estimators.

Label	Assumption	Estimator	Equation
A1	$(r \perp\!\!\!\perp y \mid x, A)$	OLS	2
A2	$(r \perp\!\!\!\perp x \mid y, A)$	OLS	5
A3	$(r \perp\!\!\!\perp x \mid y, Z)$	ML, TS, PSS, TSN, PSW	2
A4	$(y \perp\!\!\!\perp z \mid x, A)$	ML, TS, PSS, TSN	2
A5	$(r \perp\!\!\!\perp y \mid x, A, Z)$	PSW	2
A6	$(x \perp\!\!\!\perp z \mid y, A)$	ML, TS, PSS, TSN	5
A7	$(r \perp\!\!\!\perp x \mid y, A, Z)$	PSW	5

Table 3: Estimated *beta* for different samples and measures of socioeconomic status (p-values in parenthesis).

Mobility between	Full sample	No. Obs.	Restricted sample	No. Obs.	Chow test
HG_1^s and HG_1^f	0.225 (0.000)	1035	0.199 (0.000)	377	3.04 (0.017)
HG_2^s and HG_1^f	0.216 (0.000)	1092	0.185 (0.000)	388	2.185 (0.069)
HG_1^s and HG_2^f	0.183 (0.000)	1533	0.137 (0.000)	875	4.386 (0.002)
HG_2^s and HG_2^f	0.185 (0.000)	1621	0.140 (0.000)	917	4.829 (0.000)
HG_1^s and HG_3^f	0.076 (0.000)	2117	0.111 (0.000)	969	7.367 (0.000)
HG_2^s and HG_3^f	0.128 (0.000)	2311	0.130 (0.000)	1022	2.447 (0.044)

Table 4: Correcting for sample selection bias under MAR.

Methods	HG_1^s, HG_1^f	HG_2^s, HG_1^f	HG_1^s, HG_2^f	HG_2^s, HG_2^f	HG_1^s, HG_3^f	HG_2^s, HG_3^f
Full sample	0.225	0.216	0.183	0.185	0.076	0.128
p-value	0.000	0.000	0.000	0.000	0.000	0.000
Sub sample	0.199	0.185	0.137	0.140	0.111	0.130
p-value	0.000	0.000	0.000	0.000	0.000	0.000
Correction methods						
ML	0.199	0.185	0.138	0.140	0.110	0.129
p-value	0.000	0.000	0.000	0.000	0.000	0.000
TS	0.199	0.185	0.137	0.140	0.111	0.129
p-value	0.000	0.000	0.000	0.000	0.000	0.000
PSS	0.200	0.182	0.137	0.140	0.111	0.129
p-value	0.000	0.000	0.000	0.000	0.000	0.000
PSW	0.207	0.203	0.161	0.174	0.125	0.144
p-value	0.000	0.000	0.000	0.000	0.000	0.000
TSN	0.201	0.187	0.137	0.140	0.109	0.128
p-value	0.000	0.000	0.000	0.000	0.000	0.000

Table 5: Estimated b and $\sigma^2_{\tilde{x},\tilde{x}}$ for different samples and measures of socioeconomic status (p-values in parenthesis).

	b			$\sigma^2_{\tilde{x},\tilde{x}}$		
	Full	Restricted	Chow	Full	Supplemental	Equality
	Sample		Test	Sample		Test
HG_1^s and HG_1^f	0.331	0.324	0.834 (0.504)	0.118	0.137	0.906 (0.086)
HG_2^s and HG_1^f	0.231	0.258	1.915 (0.106)	0.118	0.137	0.906 (0.086)
HG_1^s and HG_2^f	0.257	0.192	1.848 (0.117)	0.110	0.137	0.845 (0.001)
HG_2^s and HG_2^f	0.193	0.160	1.601 (0.171)	0.110	0.137	0.845 (0.001)
HG_1^s and HG_3^f	0.253	0.267	111.1 (0.000)	0.295	0.137	0.445 (0.000)
HG_2^s and HG_3^f	0.293	0.252	119.3 (0.000)	0.295	0.137	0.445 (0.000)

Table 6: Correcting for sample selection bias relaxing MAR (β coefficient)

Methods	HG_1^s, HG_1^f	HG_2^s, HG_1^f	HG_1^s, HG_2^f	HG_2^s, HG_2^f	HG_1^s, HG_3^f	HG_2^s, HG_3^f
$\sigma^2_{\tilde{x},\tilde{x}}$ estimated using the Full Sample						
Full sample	0.225	0.216	0.183	0.185	0.076	0.128
p-value	0.000	0.000	0.000	0.000	0.000	0.000
$\sigma^2_{\tilde{x},\tilde{x}}$ estimated using the Supplemental Sample						
Full sample	0.207	0.215	0.161	0.180	0.158	0.273
p-value	0.000	0.000	0.000	0.000	0.000	0.000
Sub sample	0.203	0.240	0.120	0.149	0.167	0.234
p-value	0.000	0.000	0.000	0.000	0.000	0.000
Correction methods						
ML	0.210	0.246	0.124	0.154	0.153	0.234
p-value	0.000	0.000	0.000	0.000	0.000	0.000
TS	0.209	0.245	0.124	0.154	0.154	0.234
p-value	0.000	0.000	0.000	0.000	0.000	0.000
PSS	0.214	0.244	0.121	0.153	0.145	0.232
p-value	0.000	0.000	0.000	0.000	0.000	0.000
PSW	0.226	0.296	0.157	0.198	0.195	0.261
p-value	0.000	0.000	0.000	0.000	0.000	0.000
TSN	0.217	0.236	0.125	0.154	0.152	0.231
p-value	0.000	0.000	0.000	0.000	0.000	0.000

Table 7: Correcting for sample selection bias under MAR and for measurement error.

Methods	HG_1^s, HG_1^f	HG_2^s, HG_1^f	HG_1^s, HG_2^f	HG_2^s, HG_2^f	HG_1^s, HG_3^f	HG_2^s, HG_3^f
Full sample	0.232	0.222	0.177	0.179	0.196	0.331
p-value	0.000	0.000	0.000	0.000	0.000	0.000
Sub sample	0.205	0.191	0.133	0.135	0.287	0.336
p-value	0.000	0.000	0.000	0.000	0.000	0.000
Correction methods						
ML	0.205	0.191	0.133	0.135	0.284	0.333
p-value	0.000	0.000	0.000	0.000	0.000	0.000
TS	0.205	0.191	0.133	0.135	0.287	0.333
p-value	0.000	0.000	0.000	0.000	0.000	0.000
PSS	0.206	0.187	0.133	0.135	0.287	0.333
p-value	0.000	0.000	0.000	0.000	0.000	0.000
PSW	0.213	0.209	0.156	0.168	0.323	0.372
p-value	0.000	0.000	0.000	0.000	0.000	0.000
TSN	0.207	0.193	0.133	0.135	0.282	0.331
p-value	0.000	0.000	0.000	0.000	0.000	0.000

Table 8: Estimates of β using panel with different length (p-values in parenthesis).

Methods	HG_1^s, HG_1^f	HG_2^s, HG_1^f	HG_1^s, HG_2^f	HG_2^s, HG_2^f	HG_1^s, HG_3^f	HG_2^s, HG_3^f
Full sample	0.225	0.216	0.183	0.185	0.076	0.128
Restricted sample	0.199	0.185	0.137	0.140	0.111	0.130
Restricted sample	0.158	0.177	0.090	0.104	0.094	0.121
8 waves	(0.000)	(0.000)	(0.009)	(0.005)	(0.000)	(0.000)
No. obs.	298	308	612	646	670	713
Restricted sample	0.112	0.145	0.060	0.097	0.097	0.140
6 waves	(0.015)	(0.005)	(0.096)	(0.015)	(0.000)	(0.000)
No. obs.	282	297	509	542	560	602
Restricted sample	0.091	0.158	0.039	0.078	0.067	0.113
4 waves	(0.066)	(0.004)	(0.345)	(0.094)	(0.032)	(0.001)
No. obs.	270	287	418	445	458	492

Table 9: Probit models for the selection.

	Model 1		Model 2		Model 3	
Variable	Coeff	SE	Coeff	SE	Coeff	SE
born 1966-70	0.925	0.238	1.042	0.242	1.006	0.235
born 1971-75	1.694	0.238	1.757	0.242	1.773	0.235
born 1981-85	1.629	0.330	1.529	0.334	1.750	0.321
log house price	-2.174	0.290	-2.216	0.295	-2.265	0.291
Rest of the South East	-0.400	0.170	-0.471	0.173	-0.411	0.170
South West	-0.693	0.215	-0.833	0.221	-0.784	0.216
Anglia and Midlands	-1.161	0.224	-1.298	0.230	-1.282	0.226
North West	-1.100	0.247	-1.245	0.253	-1.199	0.249
Rest of the North	-0.958	0.237	-1.107	0.243	-1.081	0.239
Wales	-1.270	0.297	-1.435	0.303	-1.367	0.298
Scotland	-1.285	0.251	-1.406	0.257	-1.377	0.251
Black	-0.476	0.412	-0.568	0.414	-0.555	0.399
Indian	1.155	0.394	1.139	0.387	1.198	0.390
Pakistani, Bangladeshi	0.301	0.432	0.283	0.440	0.497	0.408
Other races	-0.426	0.460	-0.378	0.475	-0.353	0.420
Active Catholique	-0.272	0.187	-0.248	0.189	-0.214	0.186
Active Protestant	-0.194	0.135	-0.217	0.136	-0.206	0.132
Active Other religions	1.131	0.356	1.191	0.364	1.233	0.353
log HG_1^s			-0.789	0.155		
log HG_2^s					-0.412	0.128
constant	23.355	3.343	26.829	3.479	25.878	3.413
No. observations	1035		1092		1092	
LR chi2 (p-value)	440.5		258.97		258.97	
Pseudo R-squared	0.139		0.1822		0.1822	

Table 10: Tests to verify assumptions imposed by different estimators (p-values are in parenthesis).

	A1 ($r \perp\!\!\!\perp y \mid x, A$)	A2 ($r \perp\!\!\!\perp x \mid y, A$)	A3 ($r \perp\!\!\!\perp x \mid y, Z$)	A4 ($y \perp\!\!\!\perp z \mid x, A$)
HG_s^1 and HG_f^1	-2.36 (0.02)	-1.90 (0.06)	-0.97 (0.33)	2.15 (0.00)
HG_s^2 and HG_f^1	0.39 (0.69)	-2.71 (0.01)	-1.48 (0.14)	2.19 (0.00)
HG_s^1 and HG_f^2	-2.60 (0.01)	-0.63 (0.53)	0.07 (0.94)	1.98 (0.01)
HG_s^2 and HG_f^2	-0.03 (0.97)	-1.39 (0.17)	-0.64 (0.52)	1.96 (0.01)
HG_s^1 and HG_f^3	-5.13 (0.00)	16.50 (0.00)	15.56 (0.00)	3.52 (0.00)
HG_s^2 and HG_f^3	-1.39 (0.17)	17.23 (0.00)	16.28 (0.00)	3.44 (0.00)
	A5 ($r \perp\!\!\!\perp y \mid x, A, Z$)	A6 ($x \perp\!\!\!\perp z \mid y, A$)	A7 ($r \perp\!\!\!\perp x \mid y, A, Z$)	Normality
HG_s^1 and HG_f^1	-2.96 (0.00)	1.68 (0.04)	-1.48 (0.14)	16.02 (0.00)
HG_s^2 and HG_f^1	-0.41 (0.68)	1.95 (0.01)	-2.14 (0.03)	13.47 (0.00)
HG_s^1 and HG_f^2	-2.76 (0.01)	2.15 (0.00)	-0.15 (0.88)	5.51 (0.06)
HG_s^2 and HG_f^2	-0.25 (0.80)	2.36 (0.00)	-0.89 (0.37)	7.13 (0.03)
HG_s^1 and HG_f^3	-4.97 (0.00)	2.56 (0.00)	16.69 (0.00)	14.76 (0.00)
HG_s^2 and HG_f^3	-1.29 (0.20)	2.83 (0.00)	17.30 (0.00)	28.15 (0.00)