Nancy L. Schwartz Lecture

Shirtsleeves to Shirtsleeves:

The Economics of Social Mobility

Nancy L. Stokey*

University of Chicago

Delivered May 8, 1996 at the

J.L. Kellogg School of Management

Northwestern University

December 30, 1996

*I am grateful to James Heckman, Larry Jones, Robert E. Lucas, Jr., Casey Mulligan, Derek Neal, Christopher Phelan, Peter Rossi, and Paul Wolfson for helpful comments, to Yizu Yeh for research assistance, and to the National Science Foundation for financial support.

It is an honor to be here today to give this lecture, for several reasons. First, the speakers who have preceded me in this lecture series form a very distinguished group, and it is an honor to join them. Second, Northwestern University is where I began my career. It was a wonderful environment for me and I have many fond memories, so it is nice to be back. And finally, Nancy Schwartz, in addition to being an outstanding scholar, was a kind and generous friend to me when I first arrived at Northwestern, and I am glad to have this opportunity to pay tribute to her.

I. Introduction

My topic today is intergenerational mobility, and I want to begin with a brief discussion of why I think it is an important issue. I will discuss three reasons: one that is moral or philosophical and two that are pragmatic and policy oriented.

First the philosophical question. We are all brought up with the notion that the United States is a land of equal opportunity. Now the fact that we talk about equal opportunity rather than equal outcomes means that we are willing to accept at least some inequality. In fact, as a society we are evidently willing to accept a substantial degree of economic inequality. In part this is because we believe economic incentives are needed to induce people to expend effort. Society's production of goods and services depends on the willingness of individuals to go to school, to study, to acquire skills, to practice, to work, to take risks, to plan, and to innovate. These activities are difficult or risky or both, so individuals are willing to engage in them only to the extent that they expect to be rewarded. The demise of the former Soviet bloc is in part a triumph of the capitalist philosophy that economic incentives are necessary, and the evidence seems to be that such incentives necessarily lead to a substantial degree of inequality.

But we care about the sources of inequality as well as its extent, which is why we distinguish between equal opportunity and equal outcomes. To what extent is the claim that our society provides equal opportunity justified? How can we tell?

I am going to take the position that if economic success is largely unpredictable

on the basis of observed aspects of family background, than we can reasonably claim that society provides equal opportunity. There still might be significant inequality in income across individuals, due to differences in ability, hard work, luck, and so on, but I will call these unequal outcomes. On the other hand, if economic success is highly predictable on the basis of family background, then I think it is difficult to accept the claim that our society provides equal opportunity. In this case accidents of birth--unequal opportunity--are primary determinants of economic status.

Consequently, on a first pass we can judge whether there is equal opportunity by looking at parents and their children to see whether the economic success of the children is determined in large part by the success of their parents. If it is, then we can go on ask what the mechanisms are by which this happens and whether there are any changes in social or institutional arrangements that would lead to more equality of opportunity.

To think concretely about the causes and consequences of social mobility it is useful to have a couple of specific issues in mind. Two that are much in the news recently are school funding policy and immigration policy.

Local funding of public primary and secondary schools, which was the standard method of financing in this country for many years, led to highly unequal spending across school districts. In most states a substantial portion of the revenue is now collected and apportioned at the state level, resulting in much more equal expenditures per pupil across districts. Is this likely to have much effect on equality of opportunity? More generally, how do various institutional arrangements affect the educational outcomes of children from various types of families? For example, how important are funding arrangements for public schools, the availability of college scholarships and loans, the availability of publicly supported universities, and so on in determining educational outcomes for various types of children?

The number of new immigrants entering the U.S., both legal and illegal, has increased dramatically since the passage of the Immigration Act of 1965. Many of these

immigrants have substantially less education and lower earnings than the native born population, and the debate about whether our immigration policy should be tightened has become very heated. A sensible discussion of the long-run consequences of a restrictive or a liberal policy involves many issues. One of these is what the children and grandchildren of these immigrants are likely to look like in economic terms. How long will it take for the families of the current immigrants to be assimilated economically? How many generations will be required before their descendants are similar to the native born in terms of educational and economic attainment?

II. Inequality and mobility

A. Measuring inequality

We must be careful to distinguish between inequality and immobility: they are not the same. To provide a rough idea about the extent of inequality in the U.S. and how it has changed over time, Figures 1a-1c display the distribution of earnings of men age 35-40 in the U.S. for the years 1960, 1980, and 1990. Earnings, which are the sum of wage and salary income, self-employed income, and farm income, are measured in current dollars, in log form, and the data are from the Censuses for those three years.

[Insert Figure 1 about here]

The figures clearly show a substantial increase in nominal income from year to year: each of the later distributions is shifted to the right. These shifts reflect both inflation and real growth. Over the twenty year period between 1960 and 1980, the average inflation rate was about 4.4% per year and the average real (deflated) earnings for this group grew at about 2.9% per year. Between 1980 and 1990 most of the change was due to inflation: during the 1980s the average inflation rate was about 5.2% per year and average earnings for this group grew at about 1.5% per year.

The figures also show an increase in inequality from year to year: each of the

later distributions is more disperse. To get a sharper picture of how much inequality increased over this period, it is useful to compare individuals at various percentiles. This is done in Table 1. As shown there, in 1960 a man at the 75th percentile had earnings that were 134% of the median, and one at the 90th percentile had earnings that were 173% of the median. Individuals at the 25th and 10th percentiles had earnings that were 73% and 41% of the median.

[Insert Table 1 about here]

Between 1960 and 1980 individuals above the median gained a little in relative terms, while those below the median lost ground. In 1980 the earnings of individuals at the 75th and 90th percentiles rose to 136% and 178% of the median, while the earnings of individuals at the 25th and 10th percentiles fell to 65% and 38% of the median.

Between 1980 and 1990 individuals above the median enjoyed another gain in relative terms, while those below the media suffered a further loss. In 1990 the earnings of individuals at the 75th and 95th percentiles grew to 143% and 197% of the median, while those at the 25th and 10th percentiles fell slightly to 64% and 36% of the median. Overall, the most significant change was the very large increase in the relative earnings of those in the upper tail of the distribution, but all of the groups showed substantial changes.

Thus, although the 1960, 1980, and 1990 earnings distributions in Figure 1 all have roughly similar shapes, the means differ substantially and dispersion grew to some extent. The same general conclusions hold for earnings for other groups, for wage rates, and for household incomes. To adjust for growth in average income, it is convenient to measure earnings, wage rates, etc. in logarithmic form and as deviations from the cohort mean. In terms of Figure 1 this simply means relabeling the horizontal axis so the mean of each distribution is zero and choosing any convenient (but constant)

units for the scale. Throughout the rest of this lecture, whenever I talk about earnings, wage rates, income, or consumption, I will be measuring it in this way: as a percentage deviation from the cohort mean.

So measured the distributions still show a noticeable increase in inequality over the last three decades. These changes are important, they will not be the main focus of my talk. Moreover, it is not at all clear that they represent a long run trend. Earlier decades show both increases and decreases in inequality, with neither predominating. For example, during the 1940s there was a very sharp compression of the wage distribution, that was only slowly reversed over the subsequent 35 years (Goldin and Margo, 1992, Figure 1). Taking a longer view, then, earnings inequality in the U.S. has shown considerable fluctuations over time, but no obvious long run trend, and for convenience I will take it to be constant.

B. What is mobility?

To measure *immobility* it is useful to have a stylized model of the relationship between fathers and sons. Figure 2a shows an idealized earnings distribution for an age cohort of fathers. It is constructed to have a mean of zero and a variance of unity. Suppose, for illustration, that each father has exactly one son. As we just saw, the data suggest that roughly the same idealized earnings distribution represents the sons. The mean income of the sons is higher, but the shape of the distribution is about the same.

[Insert Figure 2 about here]

Assume that the earnings of each son (measured as a percentage deviation from the cohort mean) is determined in the following way. The son inherits a fixed fraction--the same fraction for everyone--of his father's relative position. In addition each son's earnings has an idiosyncratic random component, which we can interpret as ability or luck. The degree of *persistence* or immobility in the society is measured by the fraction

of the father's relative position that his son inherits. I will refer to this number, which lies between zero and one, as *beta*. Low values for beta indicate low persistence (high mobility), high values indicate high persistence (low mobility).

Suppose that each son inherits three quarters of his father's relative position, so beta = 0.75. Take fathers with earnings at the 10th, 50th, and 90th percentiles and consider their sons. (Remember that each one of our idealized fathers has one son.) We can pull these three groups of sons out and look at each group separately. These three groups are shown in Figure 2b. Notice that there is almost no overlap between the top and bottom groups. In this society if you went to two picnics with these two groups, it would be very easy to figure out which group was which.

Alternatively, suppose a son inherits only one quarter of his father's relative position, so beta = 0.25. If we do the same experiment, we get Figure 2c. In this case the earnings distributions for sons of fathers at the 10th and 90th percentiles have substantial overlap. In this society it would be hard to distinguish the two groups of sons.

If there were complete mobility, that is if a son inherited none of his father's relative status (beta = 0.0), then all groups of sons would look exactly alike. If there were complete immobility, each son would inherit exactly his father's relative position (beta = 1.0). This would be a perfect caste society.

Notice that in societies with more persistence, a father's earnings are more useful in predicting his son's earnings, and vice versa. Notice, too, that the degree of immobility--as measured by beta--can change without changing the degree of inequality--as measured by the variance of the distribution. Similarly, the degree of inequality can change without changing the degree of immobility.

C. Pitfalls in measuring social mobility

Opinion among economists about the extent of social mobility has changed dramatically in the last decade, and it is interesting to see why this is so. Ten years ago

a survey of the available empirical studies would have led to the conclusion that there was very little persistence in economic status. A number of studies, using data from the U.S. and several western European countries, all found persistence coefficients of 0.20 to 0.25. The evidence today, as we will see in just a moment, is that these estimates are far too low. Two important problems marred the early studies.[1]

The first problem was that the samples were nonrepresentative, and this biased the results. To illustrate the problem, consider the society depicted in Figure 2b, where the persistence coefficient between the earnings of fathers and sons is 0.75. Figure 3a displays earnings data from this society in a different way, with sons' (normalized, log) incomes plotted against their fathers'. In this figure the persistence coefficient is represented by the slope of the regression line.

[Insert Figure 3 about here]

Suppose that a sample from this population is constructed by taking young men, sons, who are high school graduates and who joined the army. A sample like this might be used because the data are easy to obtain. Then many of the sons with very low earnings or very high earnings do not appear in the sample: the former did not finish high school and the latter did not join the army. To make the point very clearly, suppose that none of the sons with earnings below a certain minimum or above a certain maximum appear in the sample, and that all the others do. The resulting sample appears in Figure 3b. Clearly, a regression line fit to the sample group is flatter than the regression line for the population. That is, the estimated persistence coefficient, the slope, is too low.

The example here is stark, but clearly the principle is very general: if the sample is chosen in a way that makes the sons similar to each other, then it is hard to discern the similarity between fathers and sons. The slope of the regression line fit to the sample is too low, and we underestimate the degree of persistence. This is the problem of sample selection bias.

The second problem was that fathers' earnings (or wage or income) were poorly measured. For the purpose of estimating persistence coefficients, we want to compare the *lifetime* earnings of fathers and sons. But earnings in any one year typically have a large transitory component: a spell of unemployment or poor health might reduce income, and overtime hours or moonlighting might raise it.

As before, suppose that Figure 3a represents the true relationship between the lifetime earnings of fathers and sons. If fathers' earnings are measured in only one year, the transitory component for that year scatters the points to the left and right. Figure 3c shows the sample that results and the corresponding regression line. Again, the regression line for the sample is flatter than the true regression line, and the estimated persistence coefficient is too low. One method for removing--or at least reducing--measurement error is to observe fathers' earnings for several years and take an average. Other more sophisticated techniques are also available.[2]

III. Recent estimates of persistence

Next I want to review several recent empirical studies, to give a sense of what the best current estimates suggest about the degree of intergenerational mobility. I will look at five studies, all using U.S. data. The first three look at representative U.S. samples; the last two focus on immigrant groups.

A. Evidence from U.S. family data

The first study is by Gary Solon (1992), who uses data from the Panel Study on Income Dynamics. The PSID is a nationally representative survey of about 5000 families, conducted annually since 1968. It contains data on the wage rates and earnings of individual family members, total family income (including nonlabor income), some components of consumption, and many demographic variables. In addition, information about assets (wealth) was collected in two years.[3]

This study has been running for a long enough period so that the children of the original PSID households are now adults. Consequently it is a good source of data on

the earnings and income of parents and children. A nice feature of this data is that the earnings, wages, income, etc. of the parents are self-reported values collected contemporaneously. In this respect the information should be quite accurate. Data on earnings and income collected retrospectively are likely to contain much larger reporting errors, especially if the information is collected from the children.

Solon looks at sons in the cohort born between 1951 and 1959. The 1984 earnings of the sons are related to their fathers' earnings in 1967-1971, so the sons are 25-33 years old when their earnings are measured and the fathers are 27 to 68 years old in the first year their earnings are measured. Observations where earnings are zero are dropped, and if a family has more than one son only the oldest is included. All of the regressions have around 300 observations and all control for the age of the father and the age of the son.

Table 2, which shows the persistence coefficients for a variety of regressions of son's earnings on father's earnings, clearly displays the effects of selection bias and measurement error. In the first column the sample is trimmed to include only sons who have at least twelve years of education, and the earnings of the fathers are measured in 1967. The other columns use the full sample, and each entry is calculated in the following way. For single-year measures of father's earnings, Solon estimated five regressions, using father's earnings in each of the years from 1967 to 1971 as the independent variables. For two-year measures of father's earnings, he estimated four regressions, using father's average income in 1967-68, 1968-69, etc. The coefficient in the second column here is the average of Solon's five single-year estimates, the one in the third column is the average of Solon's four two-year estimates, and so on.

[Insert Table 2 about here]

The effects of selection bias and measurement error are striking. It is clear how earlier studies could have arrived at low estimates of persistence. For the trimmed sample the persistence coefficient is 0.21, which is in the range of the early estimates.

Eliminating selection bias, by using the entire sample, raises the estimated coefficient to 0.30. Reducing the measurement error in father's earnings, by averaging reported earnings over several years, further increases the estimated coefficient. When father's income is averaged over all five years, the coefficient is 0.41.

Results from a second study confirm these conclusions. David Zimmerman (1992) conducted a similar analysis using data from the National Longitudinal Survey. The NLS, which was initiated in 1966, has data on wage rates, earnings, and the Duncan index (which is a measure of occupational status) for up to 15 years (through 1981) for about 900 father-son pairs. The average age of the fathers in 1965 was 49.7, and the average age of the sons in 1981 was 33.8. Zimmerman uses only observations where both father and son are fully employed, i.e., working at least 30 hours per week for at least 30 weeks out of the year, and if a family has more than one son only the oldest is included. All of the regressions include experience variables for fathers and sons.

Zimmerman regresses son's earnings in 1981 on single-year measures of father's earnings in 1965, 1966, 1968, and 1970, on two-year average earnings in 1965-66, 1966-67, and 1968-69, etc. The coefficients reported in Table 2 are the average of four estimates with single-year earnings, of three estimates with two-year average earnings, etc. Zimmerman finds the same pattern that Solon does, although his coefficients are somewhat higher. For single-year measures of father's earnings, the average regression coefficient is 0.40. Averaging father's income over two or more years produces steady increases in this value. For a four-year average of father's earnings over the period 1965-68 the persistence coefficient is 0.54.

Table 3 presents what I believe are the best current estimates of persistence coefficients for several different measures of status: (individual) hourly wage, (individual) earnings, (family) income, and (family) consumption. In each case the coefficients are estimated using instrumental variables, another technique for dealing with measurement error. All of the regressions include age or experience variables for both fathers and sons to control for life-cycle effects. Two sets of estimates are from the

studies by Solon and Zimmerman. The third study is by Casey Mulligan (1996), who also uses the PSID data. He uses more recent waves of that data set (up to 1989), which allows him to use several more birth cohorts. More important for the results here, he uses different instruments.

In Solon's regressions the father's status is measured in 1967 and the son's in 1984, and the instrument is father's years of education. In Zimmerman's regressions the father's status is measured in 1966 and the son's in 1981, and the instrument is the Duncan index. Mulligan measures parental variables in 1967-72 and looks at the adult children in 1984-89. The children must be born between 1951 and 1961, and there may be several adult children from one family. In addition to wages, earnings, and income, Mulligan also looks at consumption, and the persistence coefficients for the two family measures, income and consumption, include daughters. In these regressions the children may be household heads, wives of heads, or cohabitors of heads. There are about 700 observations (sons) in the wage and earnings regressions, and about 900 observations (sons and daughters) in the income and consumption regressions. His instruments are average earnings and family income in similar demographic groups, and for the consumption regression income is also used.[4]

[Insert Table 3 about here]

Although the estimates in Table 3 vary quite a bit across status measures and across studies, several patterns are clear. First, the instrumental variables (IV) estimates for persistence in earnings and in the other status measures exceed those in Table 2. Apparently, averaging five contiguous years of data on earnings still leaves substantial measurement error. Second, income and consumption apparently are more persistent than wage rates and earnings. There are several possible explanations for this. First, parents may adjust their bequests in response to a child's earnings. That is, parents may increase their bequests to low-wage children and reduce their bequests those with high wages. Assortative mating and persistence in family size may also be factors.

The estimates in Table 3 suggest that persistence coefficients for various status measures are at least 0.50 or 0.60, and perhaps even higher. This is quite a change from the early estimates of 0.20 or 0.25. Monday morning quarter-backing is easy. It is also fun, so I will indulge in saying of those who accepted the low estimates at face value that: they should have known better. Why? There were two clues. The first is Figure 2. If the true persistence coefficient were 0.25, then the sons of fathers with high earnings would be quite difficult to distinguish from the sons of those with low earnings. Moreover, in this case the persistence coefficient between grandfathers and grandsons would be $0.25^2 \approx 0.06$. Grandfathers whose earnings were twice the mean or one half the mean for their cohort would have grandsons whose earnings were, on average, only 12% above or below the mean, so the two groups of grandsons would be virtually indistinguishable from each other. Everyday observation suggests that both conclusions are wrong.[5]

B. Evidence from U.S. immigrants

The next two studies, both by George Borjas, look at immigrant groups and examine the extent to which status differentials across ethnic groups persist from one generation to the next. This work is interesting for several reasons. First, as an additional and entirely different source of information about persistence, it provides a check on the estimates in Table 3. If the children of immigrant families regress to the mean at the same rate as children of native born parents, then the persistence coefficients for ethnic groups should be the same as those for the PSID and NLS families. Persistence for ethnic groups as groups may be higher than persistence for families, however, if "neighborhood effects" that operate within ethnic groups are present. There is a great deal of controversy about the importance, if any, of such neighborhood effects in determining the educational and economic success of children. One problem that has always plagued attempts to measure them is the difficulty of identifying the relevant neighborhood. For the immigrant families in these studies, the ethnic group serves as the neighborhood. If ethnic neighborhood effects are present,

earnings differentials across ethnic groups will display more persistence than earnings differentials across families.

In his first study Borjas (1992) measures the extent to which the wage rates and occupational status measures of the children and grandchildren--both sons and daughters--of immigrants are influenced by the attainment of both their parents and their ethnic groups. Borjas looks at first and second generation Americans, i.e., those with parents or grandparents born outside the U.S. (African-Americans and native Americans are excluded.) Two data sets are used.

The General Social Surveys (GSS) consists of a series of cross sections collected between 1977 and 1989. Borjas pools information on persons age 18-64 from each wave, obtaining a total of about 6800 observations. The data include information on the occupation of each respondent and his/her father, the family's (self-reported) ethnic background, and whether the parents and/or grandparents were foreign born. There is no direct information about the wage rates, earnings, or income of the respondent or parents, so Borjas looks at mobility in the Hodge-Siegel-Rossi score of occupational prestige, which is similar to the Duncan Index. About 30 ethnic groups are represented.

The National Longitudinal Survey of Youth (NLSY) has panel data collected between 1979 and 1987. Information is available about the respondent's hourly wage rate in 1987, the birthplace (U.S. or foreign) of both parents, the (self-reported) ethnic background, and father's occupation in 1979. There is no direct information about the father's wage rate or earnings, so the analysis uses his occupational earnings, which are obtained by matching his occupation code with the average earnings for his occupation, using data from 1970 Census. The respondents are 22-29 years old when their wage rates are measured. There are about 3700 observations in the sample, and about 20 ethnic groups are represented.

Ethnic group status is the average value of the measure--occupational prestige or occupational earnings--within the ethnic group in the father's generation. Summary statistics for both data sets show wide variation across ethnic groups in both status

measures.

Before looking at the results it is useful to think about how we should interpret the coefficients from regressions that include the status of both the father and the ethnic group. If there are no ethnic neighborhood effects, then presumably the status differentials for (immigrant) ethnic *groups* will have the same degree of persistence as the differentials for (native or immigrant) *families*. If ethnic neighborhood effects are present, then ethnic groups will show *more* persistence as groups than families do.

If all variables were perfectly measured, the estimated coefficient for the parental variable would measure persistence within families, and the sum of that coefficient and the one for the ethnic group variable would measure persistence within ethnic groups. But the regression coefficient on the ethnic group variable may be positive even if neighborhood effects are absent. If father's status is measured with error, as it surely is in both of the data sets here, then the coefficient on the father's status will be biased downward, for the reasons discussed before. In this case the coefficient for the ethnic group variable will pick up some of the unmeasured effect of father's status as well as the ethnic group effect. Consequently, the sum of the coefficients on the father's status and the group's is an estimate of the persistence in status differentials displayed by ethnic groups as groups, whether or not there are neighborhood effects. *If* there are no neighborhood effects, it is also an estimate of persistence within families.

Table 4 displays the results of regressions of child's status on father's status only and on both father's and ethnic group's. All of the regressions include dummy variables for gender, for those with immigrant parent(s) (as opposed to grandparents), for those enrolled in school (in the NLSY), and for the year of the survey (in the GSS). The regressions with both father's and ethnic group's status also include age, age squared, and regional dummies.

[Insert Table 4 about here]

The persistence coefficients are modest when father's status is the only regressor: 0.20 for occupational prestige (which is measured in levels) and 0.35 for the (log) wage rate. Adding the ethnic group measure to the regression has very little effect on the coefficients on father's status and it adds a significant effect of its own. For occupational prestige, total persistence rises to 0.64, with the ethnic group measure contributing over two thirds of the total. Total persistence rises to 0.61 for the wage rate, with the ethnic group measure contributing almost half of the total.

Thus, ethnic group differentials in occupational prestige and wage rates show very high levels of persistence. The estimates, 0.64 and 0.61, are at the upper end of the range of coefficients for wage rates and earnings in Table 3. It is not surprising that persistence of economic status differentials would be as high in families of recent immigrants as in the whole population. Borjas's estimates suggest that the persistence coefficients for the ethnic groups as groups may be even a little higher.

The last study, also by Borjas (1994), looks at persistence of economic differentials by ethnic group for descendants of turn-of-the-century immigrants. During the Great Migration of 1880-1910, almost 18 million immigrants entered the U.S., an enormous influx. The U.S. population in 1880 was only 50 million, while by 1910 it had grown to 92 million (Maddison, 1982, Table B2), an increase of 42 million. Of this increase, over 40% was directly due to immigration. In addition there were indirect contributions, as the early waves of immigrants had families. Geographically, the immigrants were concentrated in the northeast and north central regions, and were more likely than native whites to live in large cities. Very few immigrants lived in the South and very few went into agriculture.

Borjas uses the 1910, 1940, and 1980 Censuses to look at these immigrants, their children, and their grandchildren, to see on average how quickly--or slowly--ethnic groups were assimilated economically. In particular, he looks at persistence in wage rate differentials for the 32 ethnic groups that constituted the bulk of the Great Migration.

First Borjas uses data for working men aged 25-64 from the 1910 Census to

construct an average earnings measure for each ethnic group. Although wage rates are not directly available, the 1910 Census does report a detailed industry/occupation code (420 categories). Borjas imputes earnings for each individual by matching these occupation codes with the occupational earnings measures for the year 1900.[6]

The 1940 Census has information about birthplace of parents, so Borjas can accurately identify second generation Americans, i.e., those with at least one parent born outside the U.S. He looks at second generation men aged 25-64 who are in the civilian sector and not in school. Both reported hourly wage and occupational hourly wage by 3-digit occupation are available. Ethnic affiliation for these men is defined by father's place of birth, unless only the mother was born outside the U.S. There is only a modest amount of marriage between ethnic groups, however, so in most cases the precise method for identifying ethnic affiliation is not crucial. Borjas looks only at those individuals in the 32 nationality groups of the Great Migration, but these are 98.3% of the second generation group.

In the 1980 Census there is no information about birthplace of parents or grandparents, so instead Borjas looks at individuals by self-reported ancestry. Of non-black natives, 80% fall into the 32 ethnic groups of the Great Migration. Borjas uses these individuals, thereby excluding only those who are black or who report no ancestry or American or an ethnic affiliation outside the 32 groups. Obviously, this sample contains most of the grandchildren of the Great Migration, but it is contaminated by the inclusion of many others. In particular, it contains second generation Americans, third generation Americans whose immigrant grandparents came after 1910, and fourth and higher generation Americans who still report an ethnic identity different from American. In addition, intermarriage across ethnic lines is more frequent among the parents of this group. As before Borjas looks at men aged 25-64 who are in the civilian sector and not in school, and focuses on reported wage rate and occupational wage by 3-digit occupation.

Table 5 displays the coefficients for regressions of (log) wage and occupational wage of the second and "third plus" generation Americans--those in the 1940 and 1980

Censuses--on the (log) occupational earnings of their ethnic group in the 1910 Census. All the regressions include age, age squared, regional dummies, and an urban dummy. There are about 11,000 observations in the 1940 sample and about 250,000 in the 1980 sample.

[Insert Table 5 about here]

For the regressions using the 1940 Census, the coefficients are 0.60 for the reported wage and 0.67 for the occupational wage. Thus, the persistence of relative status for ethnic groups in this data set is very similar to those reported in Table 4 (the sums) and also very similar to the high end of the range of values for families reported in Table 3.

For the regressions using the 1980 Census, the coefficients are 0.20 for the reported wage and 0.27 for the occupational wage. These coefficients are much smaller than the values, $0.60^2 = 0.36$ and $0.67^2 = 0.45$, that one would expect on the basis of the one-generation results. This is not surprising, since the sample from the 1980 Census is contaminated in the ways described above.

IV. Sources of persistence

A. A model of the family

Thus, the evidence is quite strong for a high degree of persistence. What are its sources? I want to think about a simple model of the family, developed and studied by Gary Becker and Nigel Tomes (1979, 1986), John Laitner (1992), Glenn Lowry (1981), Casey Mulligan (1993, 1996), Jesus Navarro (1992), and many others.

Consider individuals with finite lifetimes who care about their children. To simplify the problem I will adopt two conventions that are standard in this area, ignoring marriage and population growth. That is, I will assume that each parent, a mother, has exactly one child, a daughter, who in turn will have exactly one child when she becomes an adult, and so on. In addition I will assume that tastes are identical

within and across families.

Thus, families consisting of mother and daughter, and perhaps grandmother as well, are organized into "dynasties." There are no family ties between different dynasties and each dynasty has a constant number of living members. A period is one generation--about 30 years, and each individual lives for two or three periods.

Consider a typical 30 year period, during which the mother makes all of the decisions. The mother has an endowment of time that she can allocate to various activities, including work, leisure, and child-rearing. She also has income, the size of which depends in part on her work decision. The mother can allocate her income to various uses, including current consumption by all family members, investments in her daughter's health, education, etc., and financial bequests to her daughter. The investments include everything that will affect the daughter's earning ability later in her life, so they encompass a wide variety of specific items: prenatal care, good nutrition, health care, and educational toys, as well as expenditures on college tuition.

The motive for the mother's investment in and bequests to her daughter is altruism: the mother cares about her daughter's well-being as well as her own consumption. The daughter in turn cares about *her* daughter's well-being, so the utilities of the granddaughter and all later descendants are included indirectly. This model suggests two channels of persistence in income differentials: direct bequests of wealth and investments, like schooling, that affect the child's earnings. How important is each of these?

B. Inherited wealth

Inherited wealth is often what people think of first when they think about persistence: the Astors, the Rockefellers, and the Mellons. How important are inheritances of tangible assets for persistence in economic status? First, note that *earnings* are very persistent. Since labor income is two or three times as large as nonlabor income, the persistence in earnings is itself a very large part of the persistence in income. Table 3 suggests that differences in income are more persistent that

differences in earnings, however, so something else is going on. The data on inheritances is rather meager, but there are two studies that allow us to draw some conclusions. The first focuses on the very wealthy; the second is a representative sample of estates.

The first is a study by Paul Menchik (1979) that looks at the probate records of about a thousand Connecticut residents who died in the 1930s and 1940s leaving estates of $40,000 or more in current dollars. For about 300 of these people the death records of children were located in Connecticut, and for 200 of these cases the estate of the other parent was identified. This produced a sample of 146 sets of parents with about 200 children. Menchik compares "midparent wealth," an average of the estates of the two parents, adjusted to avoid double counting, with the child's wealth at death.

The mean midparent wealth was about $800 thousand, the median about $200 thousand. For the children the mean and median were $1,000 thousand and $150 thousand (all in 1967 dollars). Clearly, this is a very wealthy group. The coefficient in a regression of child's wealth at death on parents' wealth at death is 0.69, so evidently there is a great deal of persistence in wealth transfers in the upper tail of the distribution.

The second study is by Nigel Tomes (1981), who looks at data from a 5% random sample of estates probated in Cleveland, Ohio, in 1964-65, yielding about 650 estates. The mean estate was $12,000. Disposable personal income per capita at that time was about $2400, so the average estate was about five times per capita income. About 40% of the estates were less than $500, however, so for a large fraction of the population inheritances are negligible.

Thus, while inherited wealth may be a significant contributor to persistence in the upper tail of the income distribution, for the rest its importance is modest and for the bottom half it is apparently of little or no consequence. It is persistence in earnings that has the greater influence, because earnings are the bulk of national income and because earnings are important in all ranges of the income distribution.

C. Human capital

The other channel for persistence in economic status is the investment that parents make in their child's earning capacity. These investments in the child's human capital are guided by two factors.

The first is the expected rate of return on such investments. Empirical estimates of these rates of return are high: 17-22% for lower education, 15-16% for high school, 12-13% for college, and 7% for graduate school (Willis, 1986). Thus, investments in primary, secondary, and tertiary education all have rates of return that are significantly higher than the average rate of return on financial assets. Consequently, if parents want to increase the income of a child, investments in the education of the child are a more efficient mechanism for doing so than bequests of tangible assets. Parents will leave bequests only if the possibilities for investing in the child's education have been exhausted.

The second factor affecting the investment decision is the set of resources available to the parents. Recall that investments take two forms: direct investments of the parents' time and investments in goods and services like formal education that must be purchased in the market. The parents' decisions about both types of investment will be influenced by their wage rate(s), their nonlabor income, their skill in educating the child--which may be especially influenced by the mother's education, and the total endowment of time--which depends on whether the father is present. Both modes of investment are important: indeed, the evidence is growing that family influences in very early childhood are critical determinants of the child's success both in school and later in life. Investments in formal education are easier to measure, however, so I will focus on that decision. In particular, I will look at the relationship between family income and schooling.

The evidence is strong that family background and family income have an important effect on the probability of staying in school, at every level of education. That is, family income has a significant effect on the decisions to complete elementary school, attend and graduate from high school, attend and graduate from college, and

attend graduate school.[7] The effect at every stage is statistically significant and quantitatively important.

The theoretical model described above suggests that the influence of family income on investments in the child's education should be strong for low income families and should decline as family income rises. For low income families, investments in the child's education offer a high rate of return, so additional dollars of income will be invested there. High income families will exhaust the attractive opportunities for investing in the education of their children and invest additional savings in financial assets. The available evidence, while modest, is consistent with this prediction.

Tomes (1981) finds that family income has an important effect on the child's educational attainment only for poorer families. Specifically, he finds that among families that leave minimal size bequests to their children, family income has a significant effect on the child's educational attainment. Among families that leave substantial bequests, however, family income has no effect. Mulligan (1996) finds the same pattern in the PSID-SRC sample: parental income has a larger impact on son's educational attainment among families making small bequests. (The results change, however, if the whole PSID is used.)

The simple model described above highlights the inefficiency in educational investment created by a system where parents finance the education of their children. Since educational investments depend on the parents' income, and incomes vary across parents, education will vary across children for reasons that are unrelated to differences (in ability, for example) in the children themselves. The absence of markets for loans to invest in children is the source of the inefficiency.

There are a number of possible interpretations of this inefficiency. One is that low income families find it difficult to send their children to college: not only is the tuition cost high, but the child's entry into the labor force is delayed. Consequently, even if the investment in a college education is financially attractive, in the sense of having a high rate of return, the family may be liquidity constrained and hence unable

to make the investment. This underinvestment is inefficient, and it can be remedied at least partially by offering scholarships and loans to low income students.

Another interpretation is that many of the critical investments are made very early. The idea here is that a well prepared student from a low income family could go to college without undue financial hardship. Such students might not be able to afford to attend elite--and expensive--private institutions like Northwestern and the University of Chicago, but there are many alternatives. Illinois and all its neighbors have excellent public universities, as well as community colleges, that charge very modest tuition. In addition, many scholarships, loans, and work opportunities are available.

According to this interpretation, the main reason that so few children from low income families go to college is that so few are well prepared to do so when they finish high school. Their reading, writing, and quantitative skills are low, and given this starting point, a college education is not a particularly good investment. The cost of going to college is not the main problem, and offering more college loans may not help very much. Instead, the problem is underinvestment at a much earlier stage: children from lower income families attend lower quality schools, and therefore are less successful in acquiring the cognitive skills that are important for subsequent education. In the rest of this lecture I want to present some evidence that I think supports this point of view.

In a recent paper Derek Neal and William Johnson (1996) look at some data that reveal a great deal about the source of black-white wage inequality. I want to describe their results here because I think they are also very suggestive about the sources of wage inequality between children from high and low income families.

Using data from the National Longitudinal Survey of Youth (NLSY), Neal and Johnson look at the performance of black and white high school students on the Armed Forces Qualification Test (AFQT), a set of four tests measuring word knowledge, paragraph comprehension, mathematics knowledge, and arithmetic reasoning. Extensive research on the AFQT shows that it is a good predictor of (objective)

performance in a variety of military tasks, and that it is not racially biased. Moreover, it apparently is a measure of acquired skill, not innate ability: scores improve with age and years of education, and black-white gaps widen with age.[8]

The tests were administered to a group of 15-23 year olds in the summer of 1980. Neal and Johnson look at the 15-18 year olds, a group that had not yet started college or entered the labor force full-time. For these young adults the test measures the verbal and mathematical skills they had acquired by late high school.

The NLSY surveyed the same group again ten years later, in 1990 and 1991, and obtained measures of their hourly wage rate. At that time the blacks and Hispanics in the group had significantly lower average wage rates than the whites, and this was true for both men and women. By itself this is not surprising. Moreover, Neal and Johnson find that educational attainment explains only a small fraction, about 20%, of this wage gap.

What is surprising is that the AFQT test scores ten years earlier explain a large fraction of the wage gap: 75% for men and 100% for women. That is, the wage gap between blacks and whites in their middle to late twenties can be almost fully explained on the basis of skills they had--or hadn't--acquired by the time they had finished high school. A very modest part of this skill differential works indirectly, by influencing the decision to attend college.[9]

Thus, cognitive skills at the end of high school are a critical determinant of future wages, and policies directed toward raising those skills may be the most valuable tools for raising the earnings of children from low income families. What are such policies? An obvious solution would seem to be to provide better schools. Doing this may be harder than it seems, however.

V. School quality and student performance

The evidence on school quality and student performance is surprisingly mixed. I don't think anyone doubts that there are enormous differences among schools, and I expect that if you polled the parents of school-age children in, say, the Chicago area,

there would be considerable agreement about which schools were better than others. What is difficult, however, is to find robust connections between student performance and any of the obvious indicators of school quality. In particular, the most widely used measures, like expenditures per pupil and teacher-pupil ratios, have no power in predicting student performance.

It is clear that on average students from the public schools in the North Shore suburbs of Chicago perform much better than students from the public schools in the city. They perform better on standardized achievement tests, a higher fraction finish high school, a higher fraction start college, they have higher earnings, and so on. But the families living on the North Shore are also quite different from families in the city: in particular, they are better educated and have higher incomes. And we know that parents have a very large direct influence on the performance of their children.

Thus, while measurable indicators of school quality are higher in many suburban schools--teacher-pupil ratios, teacher salaries, and expenditure per pupil are all higher--it is impossible to attribute any of the superior performance of the students to these factors. After taking family effects into account, these measures of school quality have no discernable impact on student performance. The evidence on this point is quite striking, so it is useful to review it more closely

One common method for trying to assess school performance is to use a production function approach. The idea is to relate an "output" measure to various "inputs." Assume the output is an end-of-year score on a standardized test. (It could also be the fraction of students continuing to the next grade, the graduation rate, or the college entrance rate.) The inputs include the student's beginning-of-year test score, family characteristics like family income and parents' education, peer characteristics like average income and average education in the community, school characteristics, and teacher characteristics. Since the beginning-of-year test score is one of the inputs, this type of analysis can be used to determine how much "value added," as measured by improvement on the test, is contributed by each of the other inputs.

How can school quality be measured? That is, what school and teacher

attributes might appear in such an analysis? Instructional expenditures are about two thirds of school budgets. Clearly, these expenditures are dictated by class size and teacher salaries, and under the present system teacher salaries are determined by experience and educational attainment. Thus, the typical measures of school and teacher quality are the teacher-pupil ratio, teacher educational attainment, teacher experience, teacher salaries, and expenditures per pupil.

In a very interesting survey article, Eric Hanushek (1986) tabulates the results of 147 published studies that attempt to find a link between student performance and these measures of school and teacher quality. He finds that the vast majority of the effects reported in these studies are statistically insignificant, and among those that are statistically significant a sizable fraction have the wrong sign. Only teacher experience has some success in explaining student performance, and even there the evidence is not overwhelming. Moreover, it is unclear whether the correct interpretation is that experience is useful or that unsuccessful teachers tend to switch careers.

Does this mean that all teachers and all schools are alike? Hanushek rejects this idea emphatically and so do I. The problem is not that all teachers are equally effective, but that we are looking at the wrong measures of teacher quality. In other studies Hanushek (1971, 1992) looks at "teacher effects." This concept will be familiar to parents. If a family with school-age children moves to Evanston or Wilmette, it is not difficult for the parents to elicit opinions from other parents about which teachers in, say, the local elementary school are the better teachers. And I conjecture that there would be a fair amount of agreement among parents about who they are.

Hanushek looks for systematic data of this type. Consider first the problem of evaluating teachers within a school. Using the value added approach described above, we can relate the performance of the students to the specific teacher(s) that student had during the year. Performance must be adjusted for family variables if these vary systematically across classrooms. Since other characteristics of the school are the same for all the students, as are neighborhood effects, these can be neglected. Each teacher's performance can then be measured by looking at the average value added for the

students in his or her class(es), and teachers can be compared by comparing the average gains in the performance of their students. If the average gain is higher for some teachers than for others, year after year, we probably want to say that those individuals are better teachers. For school principals, this type of analysis is an obvious way to rate teachers.

Are teacher effects important? That is, is there systematic evidence that some teachers are better than others in improving the cognitive (and perhaps other) skills of their students? The formal evidence is weak, because there are few studies that look at the same teachers over multiple years. In an early study Hanushek (1971) found nontrivial teacher effects for second and third grade teachers: having the best teachers rather than the worst would raise the student's performance by about 0.2 grade levels per year. In anther study Hanushek (1992) finds even stronger effects for elementary school teachers. There, having a good teacher--one who is one standard deviation above the mean--rather than a bad teacher--one who is one standard deviation below the mean--adds one to one-and-a-half grade levels to the final achievement test score. Both data sets are small and nonrepresentative, however, so the evidence must be viewed as tentative.

Moreover, Hanushek has very little success relating teacher performance to measurable teacher characteristics. In the first study the only characteristic that is useful is the teacher's score on a verbal test, which might measure communication ability or general intelligence. In the second study Hanushek finds weak evidence that teachers with higher I.Q. scores are more successful and that experience is valuable. Beyond this almost nothing can be said.

Is it surprising that it is hard to predict teacher performance on the basis of easily observable attributes? I don't think so. Consider other types of service professionals: surgeons, lawyers, hair stylists, etc. No one believes that all surgeons are equally good. If we have to choose a surgeon, however, we do it not by looking at his years of education or grade point average, but by looking at his past performance. We evaluate surgeons on the basis of their success in performing operations. The same is

true for many other occupations. For the most part individuals are evaluated on the basis of performance; for teachers this has been rare. It would not be hard to do, but the current structure of the public school system does not encourage--or even allow--administrators to do it.

As noted above, comparisons across school districts are difficult because demographic factors vary. If two parallel school systems operate in the same city, however, this problem is greatly reduced. This observation suggests that comparisons of public schools and Catholic schools will provide a useful source of information about school quality. Many such comparisons have been made, but they have been hard to interpret for another reason. In Chicago or any other large city, students in the Catholic schools perform better than students in the public schools: their test scores are higher, a higher proportion go to college, and so on. Nevertheless, it is hard to distinguish how much of this is due to the superior quality of instruction in the Catholic school system and how much is simply self-selection.

Suppose that it is widely believed that the Catholic schools provide better instruction than the public schools. Then parents who are especially concerned about their children's education will put them in Catholic schools, which they can do even if they are not Catholic. Consequently, compared with the public schools the Catholic schools will have students whose parents are, on average, more concerned about the school performance of their children. But a higher level of parental concern itself raises the performance of the children, independent of school quality. Thus, the performance of Catholic school students will exceed the performance of public school students because of the effects of the parents, even if the quality of instruction is identical in both school systems. Or, if the quality of instruction is actually somewhat higher in Catholic schools, the measured differences in student performance will overstate the effect of school quality.

In a recent study Derek Neal (1997) manages to avoid the problem of selection bias by using instrumental variables, with information about the availability of Catholic schools as the instrument. Neal focuses on high schools. He finds that most Catholic

high schools--69% of the schools and 79% of the students--are in urban areas (counties with population over 250,000), and that it is in urban areas where they have a significant impact.

Specifically, Neal divides his sample, the NLSY, into four subgroups: urban vs. suburban, and white vs. black and Hispanic. His performance measures are high school graduation rates, college graduation rates, and, for the men in the sample, hourly wage rates at age 27-34. He finds that for the suburban schools there is no significant difference between the performance of Catholic and public schools students. In the urban areas, however, he does find significant differences.

For whites the gain in terms of graduation rate is modest: transferring to a Catholic school raises the probability of graduation from 75% to 85%. For blacks and Hispanics, the gain is enormous: transferring to a Catholic school raises the probability of graduation from 62% to 88%. For college graduation rates the gains are also substantial. Among urban high school graduates, attending a Catholic school raises the probability of college graduation from 31% to 42% for whites and from 16% to 30% for minorities. Finally, Catholic schooling has little effect on the wage rates of white men in the sample but does have a substantial effect on the wages of men in the minority group.

Neal also presents evidence suggesting rather strongly that the source of the effect for blacks and Hispanics in urban areas is the poor quality of the public schools available to them. For these students the Catholic schools offer an education that is comparable in quality to the one available in a typical suburban public school.

VI. Conclusions

The evidence is quite strong that in the U.S. today there is a great deal of persistence in economic status from one generation to the next. The current estimates need to be refined, but clearly there is much less social mobility than we thought just a decade ago. The persistence coefficient for relative status seems to be at least 0.50 and perhaps as high as 0.60 or 0.70. If equal opportunity is interpreted to mean that

children from rich and poor families have similar chances for economic success, then our society is very far from providing it.

Is this cause for concern? Equity and efficiency are often seen as conflicting goals and the ethical question posed in terms of a tradeoff: how much efficiency is one willing to sacrifice for more equity? The evidence I have presented here suggests that this view is misleading, at least for the inequality associated with social immobility in the earnings distribution. If earnings are based on productivities, and if productivities are influenced by schooling and other investments, then there is not necessarily a conflict between equity and efficiency: the average earnings of children from low-income families can be raised without reducing the earnings of others.

Thus, high persistence at the bottom of the earnings distribution *is* cause for concern, because it is a signal that large, high-return investment opportunities are being missed. Although genetic factors might explain a modest degree of persistence, it is difficult for me to believe that the very high degree of persistence we observe is due primarily to those factors. Instead, high persistence suggests that there is substantial underinvestment: that improvements in the quality of education for children from families at the bottom of the income distribution could enhance both equity and efficiency.

The evidence shows very clearly that cognitive skills acquired by the end of high school have a dramatic impact on subsequent earnings, and a comparison of public and Catholic high schools in urban areas strongly suggests that schools differ substantially in terms of their effectiveness in developing the cognitive skills of their students. The difficulty in finding a relationship between student performance and the school attributes that are easy to measure--expenditure per pupil, teacher salaries, and so on--should not necessarily be discouraging. Perhaps it only shows how easy it is to spend money badly. Formal evidence and casual observation both suggest that teachers vary, perhaps dramatically, in their effectiveness. Thus, school systems that identify and reward effective teaching are likely to produce students with better skills. Altering the incentive structures in the education sector might have a significant effect, then, with no

change in expenditures.

Even if additional resources are required, however, the evidence suggests that it is an investment that is well worthwhile. Equal opportunity is an ideal that has inspired many generations, and the returns from moving towards it appear to be very high.

Notes

[1] See Corcoran, et. al. (1990) and Solon (1992) for a further discussion.

[2] Because a regression line is fit by minimizing the sum of squared deviations in a vertical direction, selection bias and measurement error enter asymmetrically for fathers and sons. Specifically, using a nonrepresentative sample of fathers does not bias the slope of the regression line, although it may increase the standard error of the estimated coefficient. Similarly, measurement error for the sons does not introduce bias, although it will reduce the goodness of fit. Thus we need to worry about selection bias for the sons but not the fathers and about measurement error for the fathers but not the sons.

[3] Although the original PSID sample was chosen to be representative, attrition has introduced some biases. Studies have found that low and high income individuals--both parents and children--are more likely to leave than those with middle incomes. Thus, the sample used here to some extent under-represents fathers who are poorly educated or have low earnings.

[4] The instruments Mulligan uses (nine in total) are average 1970 earnings and family income in the same sex / race / one-digit occupation category, in the same sex / race / schooling category, and in the same two-digit industry, and average 1970 per capita personal income, per capita earnings, and earnings per worker in same county. A dummy variable for daughters is included in the regressions for income and consumption. Mulligan's consumption data is for food eaten at home and away from home, rent, and the value of owner occupied housing. The data are averaged for the available years and then weighted using weights calculated from the Consumer Expenditure Survey to estimate total household nondurable consumption. No adjustment is made for family size.

[5] Moreover, studies of persistence for occupational prestige--a Duncan index or similar measure--had found much higher coefficients. For example, see Corcoran and Jencks (1979).

[6] No occupational wage is available for farmers, who comprise about 25% of the native-born population but only a very small fraction of the immigrants. Thus, the immigrants are in effect compared with the nonagricultural native white population.

[7] See Cameron and Heckman (1996) for a survey of the evidence.

[8] It is also interesting that on this test developed to predict performance of military tasks, women have higher average scores than men!

[9] These results are consistent with several earlier studies. For example, see Corcoran and Jencks (1979) and Crouse (1979).

References

Becker, Gary, and Nigel Tomes. 1986. Human capital and the rise and fall of families, *Journal of Labor Economics*, 4: S1-39.

Becker, Gary, and Nigel Tomes. 1979. An equilibrium theory of the distribution of income and intergenerational mobility, *Journal of Political Economy*, 87: 1153-89.

Borjas, George. 1994. Long-run convergence of ethnic skill differentials: the children and grandchildren of the great migrations, *Industrial & Labor Relations Review*, 47: 553-573.

Borjas, George. 1992. Ethnic capital and intergenerational mobility, *Quarterly Journal of Economic*s, 107: 123-50.

Cameron, Stephen V., and James J. Heckman. 1996. Life cycle schooling and educational selectivity: models and evidence, manuscript, University of Chicago.

Corcoran, Mary, et. al. 1990. Effects of family and community background on economic status, *American Economic Review*, 80: 362-366.

Corcoran, Mary, and Christopher Jencks. 1979. The effects of family background. In Christopher Jencks, ed., *Who Gets Ahead?* New York: Basic Books.

Crouse, James. 1979. The effects of academic ability. In Christopher Jencks, ed., *Who Gets Ahead?* New York: Basic Books.

Goldin, Claudia, and Robert A. Margo. 1992. The great compression: the wage structure in the United States at mid-century, *Quarterly Journal of Economics*, 107: 1-34.

Hanushek, Eric A. 1992. The tradeoff between child quantity and quality, *Journal of Political Economy*, 100: 84-117.

Hanushek, Eric A. 1986. The economics of schooling: production and efficiency in public schools, *Journal of Economic Literature*, 24: 1141-77.

Hanushek, Eric A. 1972. Teacher characteristics and gains in student achievement: estimation using micro data, *American Economic Review*, 61: 280-288.

Laitner, John. 1992. Random earnings differences, lifetime liquidity constraints, and altruistic intergenerational transfers, *Journal of Economic Theory*, 58: 135-170.

Loury, Glenn C. 1981. Intergenerational transfers and the distribution of earnings. *Econometrica*, 49: 843-867.

Maddison, Angus. 1982. *Phases of Capitalist Development*, New York: Oxford University Press.

Menuchik, Paul L. 1979. Inter-generational transmission of inequality: an empirical study of wealth mobility, *Economica*, 46: 349-362.

Mulligan, Casey B. 1993. *On Intergenerational Altruism, Fertility, and the Persistence of Economic Status*. Doctoral dissertation, University of Chicago.

Mulligan, Casey B. 1996. *Parental Priorities and Economic Inequality*, Chicago: University of Chicago Press, forthcoming.

Navarro, Jesus Ignacio. 1992. *A Model of Bequests and Intergenerational Mobility*. Doctoral dissertation, University of Chicago.

Neal, Derek A. 1997. The effects of Catholic secondary schooling on educational attainment, *Journal of Labor Economics*, forthcoming.

Neal Derek A., and William R. Johnson. 1996. The role of premarket factors in black-white wage differences. *Journal of Political Economy*, 104: 869-895.

Solon, Gary R. 1992. Intergenerational income mobility in the United States. *American Economic Review* , 82: 393-408.

Tomes, Nigel. 1981. The family, inheritance, and the intergenerational transmission of inequality, *Journal of Political Economy*, 89: 928-58.

Willis, Robert J. 1986. Wage determinants: a survey and reinterpretation of human capital earnings functions. In Orley Ashenfelter and Richard Layard, eds., *Handbook of Labor Economics*, Vol. I. Amsterdam: North-Holland.

Zimmerman, David J. 1992. Regression toward mediocrity in economic stature, *American Economic Review*, 82: 409-429.

Relative position	1960 Census	1980 Census	1990 Census
10th percentile	41%	38%	36%
25th percentile	73%	65%	64%
50th percentile	100%	100%	100%
75th percentile	134%	136%	143%
90th percentile	173%	178%	197%

Table 1: Changes in inequality

Earnings relative to median at various percentiles, males age 35-40

	Trimmed sample	Single-year income	Two-year average income	Three-year average income	Four-year average income	Five-year average income
Solon (PSID)	0.21 (0.08)	0.30 (0.07)	0.36 (0.08)	0.37 (0.08)	0.39 (0.09)	0.41 (0.09)
Zimmerman (NLS)	n.a.	0.40 (0.06)	0.47 (0.07)	0.53 (0.07)	0.54 (0.08)	n.a.

Table 2: Effects of selection bias and measurement error on measured persistence of earnings (Solon 1992, p. 402 and Table 2; and Zimmerman 1992, Table 6)

	Wage (individual)	Earnings (individual)	Income (family)	Consumption (family)
Solon (PSID)	0.45 (0.10)	0.53 (0.14)	0.53 (0.12)	
Mulligan (PSID)	0.50 (0.05)	0.48 (0.07)	0.71* (0.06)	0.77* (0.05)
Zimmerman (NLS)	0.57 (0.12)	0.68 (0.13)		

Table 3: IV estimates of persistence (* includes daughters)

(Solon 1992, Table 4; Mulligan 1996, Table 8; and Zimmerman 1992, Tables 9 and 10)

	Father only	Father +	Ethnic = group	Total
Occupational prestige - GSS	0.20 (0.01)	0.18 (0.01)	0.46 (0.22)	0.64 (0.23)
Wage - NLSY	0.35 (0.03)	0.33 (0.03)	0.28 (0.10)	0.61 (0.13)

Table 4: Ethnic group influence on economic status

(Borjas 1992, Table III)

	1940 (one generation)	1980 (two generations)
reported wage	0.60 (0.14)	0.20 (0.05)
occupational wage	0.67 (0.19)	0.27 (0.09)

Table 5: Persistence in ethnic group wage differentials
(Borjas 1994, Tables 5 and 6)