Mortality, education, income and inequality among American cohorts

Angus Deaton
Christina Paxson

Research Program in Development Studies
Princeton University

June 1999

Prepared for NBER Aging group meeting, The Boulders, Carefree, Arizona, May 6–9, 1999. We gratefully acknowledge financial support from the National Institute of Aging and from the John D. and Catherine T. MacArthur Foundation for support through their network on poverty and inequality in broader perspectives. We would like to thank Anne Case, David Cutler, Finis Welch, and especially James P. Smith, for helpful comments on the first draft. This version of the paper is essentially the same at the May 1999 version, with a correction in Table 1.

ABSTRACT

People whose family income was less than $5,000 in 1980 could expect to live about 25 percent fewer years than people whose family income was greater than $50,000. We explore this finding using both individual data and a panel of aggregate birth cohorts observed from 1975 to 1995. We assume that health status is determined by social status, defined as income relative to the mean income of a reference group. When reference groups are not observed, health is a function of income whose slope (the gradient) depends on the ratio of within to between-group inequality. We derive results on how this relationship changes at different levels of aggregation. Our results on individuals show that income reduces the risk of death, and does so even controlling for education. Only some of the effect of income can plausibly be attributed to the reduction in earnings of those about to die. The panel of cohorts also shows a strongly protective effect of income, but there is evidence that cyclical increases in income may raise mortality, even when the long-run effects of income are in the opposite direction. There is no evidence that recent increases in inequality raised mortality beyond what it would otherwise have been.

Angus Deaton
221 Bendheim Hall
Princeton University
Princeton, NJ 08544

Christina Paxson
219 Bendheim Hall
Princeton University
Princeton, NJ 08544

1. Introduction

Trying to understand why mortality is so strongly related to socioeconomic status has been a major concern in demography, epidemiology, and public health for many years, and is beginning to attract the attention of economists. Data from the National Longitudinal Mortality Study (NLMS) show that people aged 25 whose family income was $5,000 or less in 1980 (and in 1980 prices) could expect to live 10 years less than those whose family income was more than $50,000, Rogot, Sorlie, Johnson, and Schmitt (1992). The concept of socioeconomic status is more widely used outside of economics than within it, and one of the issues that remains to be settled is the extent to which these income differences are caused by income, or by other factors correlated with income, education being the most obvious, see for example Fuchs (1989, 1993) and Garber (1989). Many writers believe that there is at least some direct protective effect of income, and in a recent literature much identified with the work of Richard Wilkinson (1996), it is argued that, while the first moment of income is protective, at least at the individual level, the second moment is a health hazard, so that income inequality raises mortality, if not at the individual level at least in populations or large subpopulations. Wilkinson distinguishes his hypothesis from the mechanical consequence of Jensen's inequality, that a convex relationship between mortality risk and dying means that aggregate mortality will be higher in more unequal societies. Instead he postulates that inequality itself is a health hazard and that it is less healthy for both rich and poor to live in a more unequal society. It is hardly necessary to emphasize the importance of such a link, if it indeed exists, though economists are particularly likely to mourn the loss of the Pareto criterion. The proponents of some changes, such as improvement in school quality, or raising the return on social security, make a plausible case that such changes will make everyone better-off, though some more than others. If such changes increase inequality, as almost certainly they would, the cost of lives lost would have to be offset against the economic benefits.

The original empirical support for the Wilkinson hypotheses comes from Wilkinson's cross-country comparisons within the OECD, where some measures of inequality are much more closely related to mortality levels and mortality changes than is either the level or rate of growth of national income. There are also a number of studies in the United States that find a relationship across states between income inequality and mortality, see Kaplan, Pamuk, Lynch, Cohen and Balfour (1996), Kennedy, Kawachi, and Prothrow-Stith (1996). Although this work has been challenged on a number of grounds, see for example Judge (1995), Fiscella and Franks (1997), and Mellor and Milyo (1998), and in some cases substantially modified, see Judge, Mulligan, and Benzeval (1998), McIsaac and Wilkinson (1997), and Lobmayer and Wilkinson (1999), what was initially perceived as implausible seems most recently to be commanding acceptance from careful and unbiased researchers in the field, see for example Marmot (1997) and the survey by House and Roberts (1999). While we have reservations about the robustness of many of the positive findings, our main purpose in this paper is not to try to come to judgement based on the review of the evidence—we suspect that it is too early to try to do so—but to offer some new evidence based on income, income inequality and mortality data for birth cohorts of Americans observed over the two decades from 1975 to 1995. As far as we are aware, birth cohorts have not previously been used in this context—as opposed to individual data, state data, or country data—and unlike these other sources, they offer both a cross-sectional and time series-dimension in the same data.

The next section of the paper summarizes and substantially extends a simple model first developed in an earlier paper by one of us, Deaton (1999). The model is designed to provide a framework for empirical application, and provides a way of thinking about the effects of income and income inequality in a framework in which causality runs from income to health, but where it is not absolute income that matters for health, but income relative to the average of an (unobservable) reference group. Although inequality has no direct effect on health, the fact that reference groups are not observed means that the slope of the observed relationship between health and income varies with the ratio of between- to within-group inequality. The model can be readily extended to incorporate a direct effect of inequality by making health depend on the absolute size of income differences within the reference group, but equally plausible specifications give different results so that, according to the theory, income inequality can be either protective or hazardous. We give detailed consideration to the *aggregation* of the relationship between health and income, how it can be expected to change as it is examined with different sources of data, such as individual records, averages of states or countries, or averages of birth cohorts.

Section 3 presents our empirical evidence. Data on mortality for birth cohorts are combined with data on income, income inequality, poverty, and education for 1975 through 1995 from successive Current Population Surveys (CPS) from 1976 to 1996. Deaton (1999)'s preliminary analysis used the same general approach, but worked with a shorter sample period (1981–93), different timing, a more limited range of variables, and merged the mortality and CPS data on a household basis, rather than on the much more satisfactory individual basis used here. He found that income reduced the risk of mortality, and that while the slope of the gradient was steepened when inequality was higher, there was no direct effect of inequality at the mean. In the current paper, we again document the strongly protective effects of income, and we examine how those effects vary at different points in the life-cycle. When we turn to inequality, we not only fail to find that it increases the risk of mortality, but that there is actually a *protective* effect, in apparent contradiction not only with the Wilkinson hypothesis, but with much of the theory developed in this paper. The basis for the result is the fact that when mortality was falling the most rapidly, in the late 1970s and early 1980s, years not included in the preliminary study, inequality of income was also rising rapidly. Note that because the mortality and inequality changes affected so many birth cohorts simultaneously, our results are based on more than a correlation for a few years of data. It is hard to understand why, if income inequality is so important in explaining mortality differences across states in the U.S., as well as differences between the U.S. and other developed countries, mortality should have fallen most rapidly just when inequality was rising most rapidly.

We also give a good deal of attention to the role of education, whether income is a mask for education, how income and education affect mortality in the cross-section and over time, and whether the treatment of income and education affects our results on the role of inequality. In a cross-section of birth cohorts, income and education are closely correlated so that, in order to disentangle their effects, we rely on the time-series dimension of the cohort data, supplemented by individual level data from the National Longitudinal Mortality Study. The individual level data show that both income and education are separately protective against mortality and that only some of the effect of income is removed when we attempt to allow for reverse causality from nearing death to income. In the cohort data, by contrast, income appears to *increase* the risk of mortality conditional on education, a result that we tentatively ascribe to the short-run or business-cycle effects of income on mortality.

2. Income, health, and inequality

The public health and epidemiological literatures, although richly suggestive of mechanisms, do not currently provide a precise characterization of the way in which inequality affects health. In consequence, it is difficult to know how to test the model, or how to interpret results. In this paper, we start from the framework in Deaton (1999), and show how it can be used as a basis for empirical analysis of individual level as well as aggregated data. The model is one in which the link runs from income to health though, in this context, there is nothing to stop "income" being reinterpreted as years of schooling so that, for the moment, we are not addressing the issue of whether it is income or education that is ultimately protective. However, we are not considering reverse causation, from health to income, a link that undoubtedly plays a part, although only a part, in accounting for the relationship between income and health. That income should cause health through a health Engel curve is consistent with standard health capital approaches in economics in which health is produced with health care and behavioral inputs that have to compete with leisure and other expenditures for a limited budget of time and money. But identifiable health behaviors appear to explain only about a quarter of the relationship between health and income, Marmot (1994), Lantz et al (1998), and medical care probably a good deal less, perhaps ten percent or so, Adler, Boyce, Chesney, Folkman, and Syme (1993), House and Williams (1995), so that this approach does not seem very promising. Perhaps the most promising line of investigation implicates the biochemical effects of psychosocial stress as a risk factor, linking this stress to social status, see for example Sapolsky (1993), and Cohen, Tyrell, and Smith (1991), Cohen et al, (1997) for some of the results, and Sapolsky (1996) and Adler et al (1992) for reviews.

In the spirit of this work, suppose that health status (measured positively) is an increasing function of income relative to the average income in the reference group to which the individual belongs. Write h for health, y for (the logarithm of) individual income, and z for the mean (logarithm of) income in the reference group. Once again, note that we are using income as a measure of social status, which could just as well be thought of as education, or as any other (absolute) measure of achievement. We assume that the relationship is linear, so that for $\beta > 0$,

$$E(h|y,z) = \alpha + \beta(y - z). \tag{1}$$

If we knew the reference group to which each individual belonged, we could estimate (1) directly using individual level data. However, such information is rarely available, and there can be no presumption that people relate to any clearly identifiable group, such as their neighbors, their birth cohort, or members of their same educational or occupational group. Health may be determined by status at work, status in the local (geographical) community, or status at church or in some community organization. As a result, z must be treated as an unobservable, and health conditioned only on (own) income, so that (1) becomes

$$E(h|y) = \alpha + \beta[y - E(z|y)]. \tag{2}$$

The conditional expectation on the right-hand side of (2) can be calculated from knowledge of the joint distribution of z and y. Given the linear structure, the most convenient assumptions are that z and y are joint normally distributed (recall that incomes are measured in logarithms). We write the marginal distribution of z as $N(\mu, \sigma_z^2)$, and the distribution of y conditional on z as $N(z, \sigma_\varepsilon^2)$, so that σ_ε^2 and σ_z^2 are measures of within reference group and between reference group inequality respectively.

Given the normality assumptions and the linearity of the original expectation, the expectation of health conditional on income takes the convenient form (see the Appendix for this and other useful results on the normal distribution)

$$E(h|y) = \alpha + \frac{\beta\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + \sigma_z^2}(y - \mu) \qquad (3)$$

so that the slope of the "gradient," the relationship between health and income, is a function of the ratio of within-group to between-group inequality. When reference groups are internally homogeneous relative to the disparity across groups, differences in individual income largely reflect intergroup disparity, which is irrelevant for health, so that the effect of income is attenuated. When inequality within reference groups is relatively high, individual income is a good indicator of relative income, and a good predictor of health. Another useful way of thinking about (3) is that income is an error-ridden estimate of relative income, so that the slope of the regression function is attenuated in the usual way. Note however that the convenient form in (3) depends, not only on the linearity of (1), but also the normality assumptions about reference group and within reference-group incomes.

Equation (3) does not assign to inequality any direct role in the determination of health. To the extent that changes in inequality change the relationship between intragroup and intergroup inequality, the observed relationship between health and income will become more or less steep. But even here, there is no structural effect of inequality on health. For example, if within group inequality increases, some people in each group will become less healthy, and some more healthy, but the average health of each group will not change. If intergroup inequality changes, leaving intragroup inequality the same, no one's health changes, even though the relationship between health and income changes. The gradient is driven, not by changes in health, but by changes with inequality in the relative income composition of each absolute income group.

The relative income model can be extended to incorporate a direct role for inequality. In some of the literature, there are suggestions that it the absolute income differences that matter, so that each person's health is determined (for example) by the difference between their income and the income of the best-off person in the reference group. If the best-off person is θ standard deviations from the group mean, health responds, not to the difference between y and z, but to the difference between y and $z + \theta\sigma_\varepsilon$, so that (1) becomes

$$E(h|y,z) = \alpha + \beta(y - z - \theta\sigma_\varepsilon) \tag{4}$$

Following through the argument as before, and assuming that σ_ε is orthogonal to income, gives a new version of (3),

$$E(h|y) = \alpha + \frac{\beta\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + \sigma_z^2}(y - \mu) - \beta\theta\sigma_\varepsilon \tag{5}$$

so that inequality plays a direct negative role on health as well as the indirect role originally assigned to it. This appears much closer to the kind of effect discussed by Wilkinson (1996) and in much of the related literature.

However, this account of inequality and health contains no fundamental reason why the effect must be *negative*. In the original model (1), individual health depends on relative income, so that a mean preserving spread of income within the reference group hurts some but helps others, so that average health is unaffected. In the modified model, people are hurt by their distance from the top of their reference group, so that a mean preserving spread hurts everyone except the person at the top, who is unaffected, and the healthiness of the group falls. But this is not the only way of measuring relative income effects. Instead of looking to the top, people might look to the bottom, and become healthier the further they are above the reference group floor. Just as in the upward looking model, health is determined by relative standing, and both approaches seem equally consistent with findings that people with higher social status are healthier. But the downward looking specification gives exactly the opposite effect of inequality on health. A mean preserving spread improves everyone's health except that of the poorest person, which remains unchanged, and the group becomes healthier. This model leads to exactly the same functional form as (4) and (5), but with the sign on the last term reversed. Without more content, linking health to relative economic status delivers no prediction for the direction of the effect of inequality on health.

Ambiguous although it may be, (5) is a convenient basis for empirical estimation, at least on individual level data. Health is linearly related to income, to inequality, and to an interaction term between inequality and income. For the analysis of individual level data, (5) can be used directly, for example as a probit or logit. Taking the former as an example, suppose that someone dies when h falls below some critical level h_0, then the probability of death is

$$p = \text{prob}\,(h \leq h_0) = \Phi\left(\frac{h_0 - \tilde{\alpha} - \tilde{\beta} y}{\sigma_u}\right) \tag{6}$$

where Φ is the cdf of the standard normal, σ_u is the standard deviation of health conditional on income, see (15) below, and $\tilde{\alpha}$ and $\tilde{\beta}$ are given by

$$\tilde{\alpha} = \alpha - \frac{\beta \sigma_\varepsilon^2 \mu}{\sigma_\varepsilon^2 + \sigma_z^2} - \beta \theta \sigma_\varepsilon \tag{7}$$

$$\tilde{\beta} = \frac{\beta \sigma_\varepsilon^2}{\sigma_\varepsilon^2 + \sigma_z^2}. \tag{8}$$

Note that the within-reference group inequality term σ_ε is being taken as a constant and absorbed into the intercept; even though we do not know the reference groups, or their mean incomes, we may have a proxy for reference group inequality—such as inequality over some observable group—in which case σ_ε could be included as a variable.

Although there are a number of microeconomic data sets on which (6) might be estimated, from one of which we will show results below, we are faced with the usual choice in work linking health and economics, between data sets that are rich on economic measures, but poor on health, and data sets that are strong on health, but weak on economics. One way of solving the dilemma is to merge information from multiple data sets, not at the level of the individual, but at the level of some group that is represented in more than one survey. In this paper, we group at the level of birth-cohorts observed in a particular year, and we merge mortality data from the vital registration system with income, income inequality, and education data from the Current Population Surveys. Before discussing the empirical results, we therefore need to consider the effects of cohort level aggregation on (5). We do this in a fairly general way, so as to allow other kinds of aggregation, for example by occupation, or by region. Since some of the previous work on inequality and health has used either state level or international aggregates, we need some aggregation framework if we are to compare results.

Return to equation (4) in the form

$$E(h \mid y, z) = \alpha + \beta(y - z) - \beta\theta\sigma_\varepsilon \tag{9}$$

and suppose that we observe neither y nor z, but some conditional average, denoted x. In our empirical work, x is the average (log) income of a birth cohort in a particular year, but it might just as well be an average conditioned on state or occupation. Conditional on x, y is $N(x, \sigma_c^2)$, so that if x is jointly normally distributed with both y and z, with common mean μ, then the expectation of h conditional on x is given by (again using the normal formulas)

$$E(h|x) = \alpha + \beta\left(1 - \frac{\sigma_{zx}}{\sigma_x^2}\right)(x-\mu) - \beta\theta\sigma_\varepsilon \tag{10}$$

where σ_{zx} is the covariance of x and z, mean incomes in aggregation and reference groups respectively. Once again, we have assumed that the within-reference group inequality σ_ε is orthogonal to the aggregation group income. Equation (10) is analogous to (5) and both come from projection of the fundamental behavioral relationship (1) on to different variables. They are identical if each person is her or his own aggregation group, so that individual income y and group income z coincide. But more generally, (10) and (5) are different; the bias to the slope now depends, not on the ratio of within- to between reference group variance, but (negatively) on σ_{zx}/σ_x^2, the slope of the regression of reference group on aggregation group income. When aggregation group income moves one for one with reference group income, as when the reference and aggregation groups coincide, the relationship between health and group income is lost. To get an unattenuated slope, we need to select aggregation groups whose average incomes are uncorrelated with reference group incomes.

Some examples clarify the implications of (10). While in principle it is possible for health to be negatively related to income at the aggregation group level, this seems unlikely in practice. Consider a component model of income, written as

$$y = \mu + \theta_a + \eta_t + \gamma_g + \epsilon \tag{11}$$

where, as always, μ is the grand mean, and the other components are zero-mean random terms associated with the effects of age (a), time (t), occupational group (g), and an individual idiosyncratic term. In case A, suppose that the reference group is people of the same age, in the same occupation, today. In case B, the reference group is all members of the profession today, irrespective of age. In both cases, the aggregation groups are birth cohorts, people of the same age today, so that the aggregation group income is $\mu + \theta_a + \eta_t$. In case A, the reference income is the sum of the first four terms on the right-hand side of (11), so that $y - z$ is simply the idiosyncratic residual ϵ which, by assumption, is orthogonal to z, so that the coefficient on x in (10) is zero. In terms of that equation, as is easily checked, we have

$$\operatorname{cov}(x,z) = \sigma_{zx} = \sigma_a^2 + \sigma_t^2 = \sigma_x^2 = \operatorname{var} x \tag{12}$$

so that x has no effect in the regression and aggregation-group income does not predict aggregation-group health. Even though income predicts health in the micro data, only inequality predicts health in the aggregate. Case B is different. Because the age component is included in aggregation-group income, but not in reference group income, the covariance of x and z is only σ_t^2, so that with the variance of x unchanged, (10) becomes

$$E(h|x) = \alpha + \beta\left(\frac{\sigma_a^2}{\sigma_t^2 + \sigma_a^2}\right)(x - \mu) - \beta\theta\sigma_\varepsilon \tag{13}$$

so that aggregation-group income matters for aggregation-group health, albeit with an attenuated effect. In this case, there will be attenuation except in the implausible case where $\sigma_t^2 = 0$ and there are no common aggregate shocks to income.

In general, if the reference groups are more finely defined than the aggregation groups, for example economists of the same age observed at the same time versus an age cohort, aggregation will annihilate the relationship between income and health in the aggregated data. If neither group is more finely defined—economists of all ages versus birth cohorts—or if the reference group is less finely defined than the aggregation group—for example, the whole population versus birth cohorts—income of the aggregation groups will predict average health in the aggregation groups. Note that if the reference group is the whole population, so that reference group income is the population mean, the covariance of x and z is zero so that, by (10), there is no attenuation, and the microeconomic relationship carries through directly to the aggregate relationship. Finally, if the aggregation groups are individuals, so that $x = y$, it is easily checked that (10) reduces to the attenuated micro relationship (5). This kind of analysis seems to capture some (although not all) of the discussion in Wilkinson (1997) who argues that looking across small geographical areas, income will be more important than inequality, while the opposite will be true in comparisons over aggregates for large areas, such a states or regions.

In applications such as the present, where health is measured only to the extent that it does or does not fall beneath a threshold and results in death, the parameters in (5) or (10) can only be estimated up to scale, where the scale is the standard deviation of health around the two regression lines (5) and (10). As was the case for the slopes, the residual variances come from standard results on the normal distribution. If we write σ_h^2 for the variance of health conditional on both y and z in (1), then the residual variance of the regression function (10) is

$$\sigma^2 = \sigma_h^2 + \beta^2 \left(\sigma_\varepsilon^2 - \frac{(\sigma_x^2 - \sigma_{zx})^2}{\sigma_x^2} \right). \tag{14}$$

When the model is estimated on group-aggregated data, here on age cohorts, the slope of the income relationship should be the ratio of $\beta(1 - \sigma_{zx}/\sigma_x^2)$ to the square root of (14). In the micro data, corresponding to (4), (14) also holds but with $x = y$, so that rearranging, we have

$$\sigma_u^2 = \sigma_h^2 + \frac{\beta^2 \sigma_\varepsilon^2 \sigma_z^2}{\sigma_z^2 + \sigma_\varepsilon^2} \tag{15}$$

Unfortunately, without knowledge of the variances, there is no general inequality that holds between (14) and (15), nor indeed between the coefficients on income in the micro and macro regressions, (5) and (10). The theory does not deliver any general basis for comparing estimated gradients at different levels of aggregation.

Nevertheless, there are two polar cases where the results are straightforward and which are worth keeping in mind when interpreting the results. These are (a) when the reference groups are small, and are contained within the aggregation groups, and (b) when reference groups are universal, so that reference income z is the same for all people and equal to mean income. Case (b) would arise if health is a function of relative income, but all people are members of the same reference group, or if health is simply a function of absolute income rather than relative income. In these two cases the gradients from micro and macro data can be compared. In case (a), with individual data, the gradient will be the ratio of (8) to the square root of (15), an attenuated but still positive effect. In the same case, but with aggregation-group data, there will be no relationship between aggregation group income and mortality, so that only inequality matters. In case (b), when z is common to everyone, σ_z^2 and σ_{zx} are both zero. In the individual-level data, by substituting $\sigma_z^2=0$, the ratio of (8) to the square root of (15) reduces to

$$\beta / \sigma_h. \tag{16}$$

By contrast, in the aggregation-group data, by substituting $\sigma_{zx} = 0$, we get, for the ratio of the slope in (10) to the square root of (14),

$$\frac{\beta}{\sigma_h \sqrt{1 + \beta^2 \sigma_c^2 / \sigma_h^2}} \tag{17}$$

where σ_c^2 is the within aggregation-group variance $\sigma_\varepsilon^2 - \sigma_x^2$. Probits run at the aggregation-group level will therefore be attenuated compared with those run on the individual level data. In our econometric analysis below, β / σ_h from the micro data is around –0.3, while the mean within-cohort variance of log income is about 0.6, so that the factor in the square root of the denominator of (17) is 1.054, and the two sets of estimates should be close.

What can be said about the relationship between inequality and the probability of dying at the cohort or other group-aggregate level? The formulas are sufficiently complex to permit a wide variety of results. However, if all the variances (and covariances) were to increase in proportion, which is one way in which inequality might increase, there would be two distinct effects. The first acts through the last term on the right-hand side of (10); when people are less healthy when they are lower in the reference group, average health declines with inequality, and this effect is not altered by the aggregation. The second effect can be seen by noting that the slope in (10) is unaffected by a proportionate change in variances, while the variance in (14) will move with the other variances. This is an aggregation effect which, in the aggregate data, acts to attenuate further the estimate of slope but, provided the probability of death is less than a half, will raise the probability of death. When mortality risk is small, we are on the convex portion of the relationship between the probability of death and (log) income, so that Jensen's inequality makes the aggregate probability an increasing function of the variance. Of course, this effect could be weakened (or strengthened) by relaxing the original linearity assumption between health and income. The possibility remains that changes in inequality, by changing the ratio of σ_{zx} to σ_x^2, will alter the slope itself, but the formulas are complicated enough to prevent further conjecture.

3. Data and results

Our empirical analysis is based on merging data on all-cause mortality for the U.S. as a whole with data on incomes from successive years 1975 through 1995 from the 1976 to 1996 Current Population Survey (CPS). (Note that the March CPS collects data on the previous year's income) Merging is at the level of birth cohorts by sex so that, for example, we will be relating the fraction of men or women born in year b who died in year t, to the incomes of that same cohort in year t (or possibly earlier.) The data on mortality are taken from the Berkeley Mortality Database (BMD); we use the 1 x 1 table, i.e. fractions dying by sex, by single year of age and in each year from 1900, see Wilmoth (1999). The CPS data are used to attribute to each individual a family income per adult equivalent, y say, and we then calculate various characteristics of the distribution of y over individuals for each birth cohort. In particular, we work with the cohort average of $\ln y$, with the variance of the logarithm of y within the cohort, with the gini coefficient, and with the proportion of people below the poverty line. Household income is the CPS measure of total household income and adult equivalents are measured as the number of adults plus half the number of children, defined as those under 18 years of age. Note that, although y is family income per adult equivalent, identical values of which are assigned to all members of each household, the mean and other statistics will generally differ by sex if, for example, women on average live in families with lower income. We also use data on education from the CPS, calculating years of education for each person, as well as dummies for various attainment levels so that, at the cohort level, we have data on average years of education as well as on the fraction who have graduated from high school, college, and so on. After 1991, the CPS reports education in bracketed intervals rather than in years. We use means within brackets from 1991 to attribute years of education in the later surveys. With an age restriction from 25 to 85, and with CPS data from 1975 through 1995, our data set consists of 1,281 age-year cohort averages.

Given that we have drawn a direct link from the theory at the individual level to its implications for birth cohorts, estimation at the cohort level is a viable alternative to estimation using individual records. Cohort data also have some distinct advantages of their own, as well as some disadvantages. On the positive side, because deaths are rare events, very large numbers of individuals need to be sampled to make individual data useful. By contrast, the cohort data use the data on all deaths in the U.S., eliminating the need for sampling. Cohort data also overcome one of the major difficulties in linking mortality to socio-economic status, which is the lack of individual level surveys that record adequate economic data together with information on deaths. In this paper, we link mortality information from the vital registration system with the CPS, which is probably the best source of data on incomes and education. It would be possible to extend this principle further, bringing in information on risk factors, such as smoking, drinking, and obesity, from other sources, such as the Behavioral Risk Factor Surveillance System, something we plan to do in our future work.

There are also disadvantages of cohort data, primarily the impossibility of separating genuine individual effects from those that result from aggregation. Using individual data that include community characteristics, we could in principle test directly for effects of community means or community inequality, as well as controlling for individual income. With aggregate data, we have to face the problem of identifying the direct effects from the aggregation effects, and such identification will not always be possible. One of our (subsidiary) aims in this paper is to assess the cohort approach, and to test whether it yields results that are similar to those from individual record data. If so, we will have more confidence in using it in situations where individual data cannot be used.

Figure 1 shows the log odds of dying by age for a selection of birth cohorts born from 1870 through to 1970. Although we can use only a fraction of this information in the analysis, the long-run information is important for interpreting recent events. One immediate feature of the patterns in Figure 1 is the age-profile of mortality. The risk of death is high immediately after birth, falls to low levels in the mid-teens, and then rises with age thereafter. After about age 30, the log odds of death is approximately linear in age, see also Elo and Preston (1996), albeit with a time-dependent slope. Mortality has also been falling over time, so that the age-profiles for the later born cohorts are below the profiles for those born earlier. But these two obvious features are far from exhausting these data; the log odds are far from being completely "explained" by a sum of age and cohort effects. For example, there are clear traces of specific events, such as the 1919 influenza pandemic on the cohort born in 1890. Note also that the proportional reduction in mortality is larger at younger ages—which in the limit must be true because everyone dies eventually—but that the reductions differ by sex; note for example the large reduction in female mortality during child-bearing years. The reduction in mortality among young males is much less than among young females, and for recent cohorts, there has been little or no reduction in the mortality of males in their early 20s. As a result, for the cohort born in 1973, the ratio of mortality rates for males to females in 1995, at age 22, is 3.25 to 1. For recent cohorts male mortality rates fall with age from the early 20s to the mid 30s. The causes of these deaths—violence, accidents, and AIDS—are quite different from those at higher ages, and it is implausible that a single explanation in terms of income and income inequality will do for both.

Figures 2 show the mortality data in a way that is more immediately relevant to the task at hand. The graphs show estimated year and cohort effects in the log odds of mortality for males and females separately. These figures use only the birth cohorts that are observed in at least one year between 1975 and 1995 inclusive, and which can therefore be matched to the CPS data. The left-hand figure are the year effects estimated from a regression of the log odds of dying on a set of year and age dummies, one for each year and each year of age. The right-hand figure is obtained in the same way, but with cohort (date of birth) effects replacing year effects. The left-hand figure corresponds to a fitted model in which the age-profile remains constant, but drifts down with time, so that all people alive at any given date benefit from that year's reduction in mortality. It shows that the rate of mortality decline was relatively rapid from 1975 through to the early 1980s, but has been a good deal slower since then. The timing of the increase in income inequality is not identical to the timing of the slowdown in mortality decline; the largest increase took place in the late 1970s and early 1980s, and the increase since then has been more modest. The cohort effects in the right-hand figure do not tell a picture of continuing progress. In particular, males born since 1950 show sharply higher mortality rates than those born in the 1940s. Note that those born most recently are only observed in the early period of their lives, so that these graphs give undue weight to the recent mortality experience among young adults, where deaths are due to violence, accidents and AIDS, and which do not share the long-run downward trend.

Table 1 presents a first set of results focusing on the relationship between income and mortality. The dependent variable is the log odds of dying while the independent variables are dummy variables for each age from 25 to 85 and the logarithm of income per adult equivalent. Because the age effects are removed, identification comes from the relationship between cohort (and interaction) components of mortality and income less its age profile. One hypothesis is that all of the trend reduction in mortality is potentially attributable to income, but the historical and international evidence suggest that this is not the case, see in particular Preston (1976). In consequence, and in order to avoid the risk of spurious correlation between time trends, we present results with and without the inclusion of a time trend. In the OLS regressions in the top panel of the table, the coefficient on the logarithm of family income per equivalent is –0.559 for men, and –0.528 for women; with the introduction of a time trend, these numbers become –0.281 and –0.125 respectively. Since the log odds is approximately the log probability when the probability is small, these estimates can be thought of as elasticities. If the coefficient were –0.5, a fourfold increase in income would cut the risk of death in half.

In order to compare the cohort results with those from individual level data, we estimated logits for the probability of dying using the National Longitudinal Mortality Study (NLMS). This is a survey of individuals originally sampled in CPS surveys and in the census around 1980 into which death certificates have been retrospectively merged. For all individuals from age 25 to 85, we constructed an indicator of whether the individual had died within 365 days of the interview, and estimated a logit model in which the independent variables are a set of age dummies, one for each age, and the logarithm of real family income per adult equivalent. (The public use version of the NLMS provides only real income classes, not actual income, and we have used the classes to construct a (rough) measure of the logarithm of total family income. We use the 1981 CPS to calculate mean log income by age and sex within each of the NLMS income brackets, using the results to impute log income to each person in the NLMS within their reported income bracket. In practice, this gives very similar results to the simpler method of setting income to the middle of the bracket. Adult equivalents were computed by linking individuals to households and counting the numbers of adults and children.) Figure 3 shows the one year log odds of mortality from the NLMS, computed from the fractions in the sample who died at each age, together with the age-profile of the log odds of mortality from the BMD for 1982. NLMS mortality, which excludes deaths among the institutionalized population, is somewhat lower among the elderly, though the major effect is presumably beyond age 85, and is noisier, especially at young ages. But the two data sets are clearly measuring the same thing, so that the NLMS results are a useful comparison. The coefficient on the (log) income term in the logit is –0.352 for men (161,472 observations, t=10.8) and –0.262 for women (177,953 observations, t=6.3), close to the estimates in Elo and Preston (1996) using the same data but a different specification. As we have seen, there are a number of reasons why these numbers should *not* be the same as those in Table 1. In addition to the aggregation issues, note that the NLMS records only a single estimate of family income so that, if permanent income is a better predictor of mortality than current income, or if income is measured with error, the NLMS estimate will be attenuated relative to the cohort based estimates to the extent that averaging over cohorts limits measurement error or proxies permanent income. Note also that the NLMS relates only to income in a few years around 1980, so that the two sets of calculations cover quite different periods. Given this and the aggregation differences, we find the estimates surprisingly similar!

Table 1 shows a number of other results, including breakdowns by age, and using two alternative instrumental variable strategies to estimate the income effect. In the four OLS columns, with results for four age groups, we obtain the now standard results that income matters most for health in late middle-age, in the 40–54 age range for women and 55-69 age range for men, and the effect diminishes with age. Even so, income is still protective for the oldest group of both men and women, though the result vanishes (and indeed income becomes hazardous over age 55) if time trends are included in the regressions. Perhaps most striking is the *positive* effect of income on mortality among young men; the effect is large and negative among young women. This result is associated with the recent increase in mortality rates among young men, which appears to be greater among cohorts with higher average incomes. Plausible arguments can be made for a positive association between income and the specific causes of mortality in this age group—AIDS, violence, and accidents—and there is good evidence from pooled state and time series evidence that gives very similar results, Ruhm (1999).

The way in which the protective effect of income varies with age can also be assessed directly, by entering age into the regressions not only in levels, but also interacted with income. We have done this for both the cohort and NLMS data, in the latter by running a logit on age dummies, the log of income per adult equivalent, and on age dummies interacted with the income term. The results are shown in graphical form in Figure 4, for males in the left-hand panel and for females in the right-hand panel. These graphs are noisy, particularly for the NLMS, which is to be expected given the very large number of age coefficients in the regressions. Nevertheless, for women, we see the familiar pattern whereby protection increases with age to around age 45 and decreases thereafter. For men, the two graphs are not the same. In particular, the cohort data again show a *positive* association of mortality and income among young men, a phenomenon of which there is no trace in the NLMS data. The obvious explanation here is the fact that the NLMS measures mortality in the early 1980s, before the youth mortality phenomenon had become so pronounced. More disappointing is the failure of the NLMS to show any other age pattern; perhaps we are asking too much of single year mortality data.

The instrumental variable estimates in Table 1 are motivated by an attempt to replace current income with a longer-term, or permanent measure. One strategy is to instrument income with years of education, on the assumption that education does not affect health directly—a controversial supposition that we shall investigate more directly below. Another is to use cohort dummies as instruments; this would be correct if we wanted to measure lifetime resources, which remain constant over time at the cohort level. These instrumentation schemes are likely to work better for people of working ages than for the elderly, whose current income is less well predicted by their education, and whose lifetime income will not be adequately captured by averaging over the maximum of 20 years in our sample. Nevertheless, the estimated income effects move as would be expected if long-run income is a better predictor of health than current income. In the top half of the table, instrumentation of income by years of education moves the coefficient for men from –0.56 to –0.92, and for women from –0.53 to –0.84. Instrumentation also removes much of the difference between the estimates with and without time trends, although income is still hazardous for those in the 55-69 age range when time trends are includes. Most remarkably, when years of education (but not education and cohort dummies) is used as an instrument, the protective effect of income for young men is restored. If this result is accepted— and there are reasons for not doing so, including the possible direct role of education, and the different result when cohort dummies are added to the set of instruments—it indicates that mortality among young men is procyclical, again in line with Ruhm's result that unemployment is good for health.

Figure 5 shows one aspect of the partial correlation between lny and mortality. These graphs show the residuals of regressions of mortality, years of education, and lny on a set of age dummies, averaged over birth-cohorts in the top two panels, and over years in the bottom two panels. The very clear inverse variation is clear for both men and women, as is the fact that education is a very good predictor of income at the cohort level (though less so on the time series) which helps in instrumentation, but which in our later tests, will also make it difficult to separate out the roles of income and education on the cohort data. (Note that the regressions contain more information than shown in the Figures, since the observations are cohort-year pairs, with interactions, not cohort or year averages.)

Table 2 repeats the top panel of Table 1, but with the addition of two within-cohort inequality measures, the variance of log income and the gini coefficient. (These regressions were also run with time-trends, but the results are not much affected.) The introduction of either the variance of log income or the gini coefficient produces an estimated *protective* effect of inequality on health. To see the size of the coefficients, we note that, averaged over all cohorts, log odds of male mortality declined by 0.20 from 1975 to 1995. The mean of the logarithm of income per adult income rose by 0.20, and its variance by 0.196 Using the coefficients for all age groups combined, then the rise in income equal would be predicted to reduce the log odds by about half of the actual decline, 0.10 (OLS) and 0.12 (IV). The rise in the variance would lead to a decline in the log odds of death by 0.03 (OLS) and (an absurd) 0.12 (IV). The changes in mortality, income, and inequality were much the same for females as for males, but the coefficients are different. The predicted effect of the rise in income is a decline in the log odds of female mortality of 0.08 (OLS) and 0.07 (IV), while the predicted declines from the increase in inequality are 0.03 (OLS) and 0.11 (IV). In the four OLS regressions (men versus women, variance of logs versus gini), the introduction of the inequality measure has very little effect on the income coefficient, which remains protective as before, but in all cases the inequality measure appears with a significant negative coefficient. The protective effect of inequality is about the same for males and females, and is a good deal larger for people from 35–59 than for those aged 60 and over. For the elderly, the estimated coefficient on the gini is insignificantly or barely significantly different from zero. While the theory allows for the possibility that inequality is protective—for example, if people look to the bottom of their reference group in assessing their health, or if the covariance of mean cohort and mean reference group income is large enough—the result remains implausible.

Figure 6 provides some insight into the source of these results, as well as an explanation for the difference with Deaton (1999), whose similar regressions yielded essentially no effect. The top two panels of the figure show the raw income inequality data, on the left for the gini coefficient, and on the right for the variance of logarithms. Each of these has a "male" and "female" variant because we use the CPS data to assign (the same) family income per equivalent to each person in the household, so that differences in inequality by sex reflect the distribution of men and women across households. As far as trends are concerned, there is no important difference by sex. Note also that the inequality measures are not the usual inequality measures for the whole country, but the average over all the cohorts in a given year of the inequality measure within the cohort. Such measures exclude the contribution of between cohort inequality to the national aggregate. Even so, the trends are similar to the national trends. Family income inequality rises rapidly from the late 1970s to the mid 1980s, with relatively little change thereafter. (The sharp increase in the gini in 1995 is not mirrored in the variance of logs, because it is associated with inequality at the very top of the distribution, and is probably distorted by topcoding effects and by other changes in interview protocols.)

The bottom panels of the figure are constructed by calculating the residuals of the regression of the log odds of death on age effects and on the mean of log income per equivalent, averaging by year, and plotting them against the similarly averaged residuals of the inequality measure on age effects and the mean of log income per equivalent. Without the averaging, the regression of the mortality residuals on the inequality residuals would reproduce the coefficients in Table 2, and the averaging is used only to produce uncluttered graphs. The negative (partial) relationship between inequality and mortality is transparent in both figures. When inequality was low, mortality was high (and falling), and when inequality was high, mortality was low. The timing of mortality change and income inequality in the U.S. is not supportive of the hypothesis that inequality increases the risk of death in the aggregate. The estimated protective effect of inequality is reduced if time trends are included in the regression, but neither the sign nor the statistical significance is altered. Nor are the results affected by introducing one or two period lags between inequality and mortality.

Table 2 also shows the consequences of instrumenting both inequality and income, on the same grounds as before, that it is possibly the long-term experience of high income and high inequality, not their year to year variations, that conditions mortality. Also as before, we use education as instruments, not only the mean years of education in the cohort, but also the fractions of the cohort with various educational attainment levels. The inclusion of the latter captures the distribution of educational attainment within the cohort, and generates excellent instruments for income inequality. Generally, the results are what might be expected, that moving to a long term basis reduces attenuation and makes the estimates absolutely larger, for both mean income and inequality. There is one exception, for men aged 39 to 59, where the protective effect of inequality is reversed, so that we get the (originally expected) positive coefficient, though for neither the gini nor the variance of logs is the estimate significantly different from zero. (One line of investigation that needs to be pursued with different data is the extent to which these results reflect mortality among young men, and its relationship to the business cycle on the one hand, and with income in the cross-section through AIDS, an effect that has almost certainly changed sign over time.)

Table 3 investigates another possibility that is often raised in the literature, that the effects of inequality might be a mistaken attribution of the effects of poverty. The first two rows of the table show the OLS and IV results for men and women, first including only the fraction of people below the official US poverty line, and in the next row, including both the fraction poor and the variance of log income. These results merely serve to deepen the puzzle. Poverty, like inequality, is estimated to be protective of health, and when both poverty and inequality are included, both are separately protective.

The last row of the table investigates whether the slope of the gradient between income and mortality is affected by mortality. As we saw in the theoretical development, except when we observe individual and reference group income, the effect of income on health is a function of inequality, and is predicted to increase with general increases in inequality in the individual level data if such changes increase within-group inequality more than between-group inequality. In the cohort data, the same effect is caused by an increase in the variance of aggregation-group income relative to reference-group income, which is itself a plausible consequence of an increase in within-reference group inequality. (For example, see (11), and take $z = \mu + \eta_t$ and $x = \mu + \theta_a + \eta_t$, so that increases in the variance σ_a^2 will simultaneously increase the slope of the gradient as well as the ratio of within-reference group to between reference group inequality.) The interaction terms in the third panel of the table are estimated to be negative, so that the gradient of mortality with income is steeper when inequality is larger. This is again a plausible result, but it does not remove the implausible (and significant) protective effect of inequality at the mean, see the derivatives at the bottom of the table.

In the results so far, we have adopted the position that the underlying determinant of health is income rather than education, and have used education to instrument income and income inequality. But this is a controversial position; many would argue that it is not income, but education—or at least some personal attribute that is strongly related to education—that is the ultimate determinant of health. It is also possible, in the immediate context, that our misspecification of the role of education is responsible for our unexpected results on the role of inequality. Americans have become more educated on average, and the rate of return to education has risen, so that changes in income inequality are in part due to changes in the relationship between income and education.

Before looking at the cohort data, it is worth using the individual level data in the NLMS to give another assessment of the income versus education, and one that is not affected by the time-series effects that are potentially important in the cohort data. Table 4 shows results from logit estimation of the probability of death on men and women in the NLMS within a year of interview.

Age effects by single years of age were included but are not shown. In the first row are the results already discussed, on the effects of the logarithm of income per adult equivalent on the log odds of dying. The second row shows what happens if we replace income by years of education; like income, education has a strong and significant effect on mortality; an additional five years of education, from say high school to a masters, reduced the probability of death by around 20 percent for men and about 25 percent for women. In the final row of the table, the logits are run with both education and income. For men, the income coefficient is reduced hardly at all, while the coefficient on education is reduced threefold, and is no longer significantly different from zero. According to these estimates, to a first approximation, it is income, not education, that is protective of health. These results are not replicated for women whose combined regression shows effects of both income and education separately, each protective, and each with a coefficient somewhat smaller than when they are included alone.

Income is much less well predicted by education for women than for men, but this does not explain why it is income, not education, that plays the dominant role in male mortality. Although others in the literature have found that income drives out education—see Lantz et al (1998) who use the (much smaller) American Changing Lives Survey, which also allows controls for behavioral factors—we find the result quite surprising. Even if it is ultimately income that matters, it is astonishing that a single observation of a year's income—with all the usual measurement error—should predict mortality better than a longer term measure as predicted by education, in which case both income and education should show up in the reduced form regression. The obvious possibility is that there is causality running from health to income for people about to die. There is some evidence for this in the results reported by Elo and Preston (1996), who also find marked reductions in the education effects on five-year mortality when income is introduced, with reductions larger for men. However, even after allowing for income—and for a range of other covariates, but with age effects restricted to entering linearly— male education is still significant, and twice as large as in Table 4.

Tables 5 (men) and 6 (women) report results from the NLMS is the same format as in Table 4, but using deaths in periods at various lengths after the interview; in the first panel for death in the first year, as in Table 4, in the second panel for the second year, the third panel between two to five years after the interview, and in the last panel, for deaths from five to nine years after the interview. (For the results in the second, third, and fourth panels, the logits are estimated only over the group of individuals who survived to the beginning of the period, so that the samples become successively smaller across the panels.) Although moving forward in time will not eliminate the effect of prospective death on income—some conditions will produce low income for many years prior to death—it should certainly reduce the influence of reverse causality. And indeed, the results in the table are supportive of such an interpretation. For men under 60, in the last two lines of the top panel, the effects of education on the log odds hold fairly steady, but the effects of income are reduced. Some such effect is also to be expected from the increasing irrelevance of an increasingly remote measure of income, so it is not clear how much of the reduction should be attributed to reverse causality. But for women, in Table 6, where the reverse causality is weaker, the initial income estimates are lower, and are less affected as the mortality window is moved forward. On this evidence, while the initial male estimate of –0.4 is probably too large (in absolute value), a case can be made for defending an estimate of around –0.3. Even for the elderly, and even five to nine years after interview, income exerts a protective effect against mortality, and the effect is not removed by controlling for education.

The cohort data are perhaps less well suited to investigating the question of whether it is income or education that matters for health. As we saw in Figure 5, the cohort average of the mean of the logarithm of income per equivalent adult is closely related to the mean years of education, even after removing age effects from both. As a result, attempts to include both variables in the cohort regressions lead to a good deal of instability in the results. Nevertheless, the cohort data allow us to investigate the possibility that there are dynamic effects of income on mortality, with differences in short and long-run responses. Table 7 shows the results of trying to investigate education versus income, and the effects on the estimates of variance. The first two rows, labeled "regression 1" shows the results of OLS regressions on age effects and on mean income and mean years of education. For both men and women, and whether or not time trends are included, income is either hazardous (or insignificant), and education is protective. The absence of a protective role for income is in sharp contradiction to the cross-sectional results from the NLMS, and for males, that education drives out income is exactly the opposite of the NLMS result. That income might actually be harmful once education has been controlled for has been argued by Fuchs (1974, 1993) and by Garber (1989), but the studies cited either do not support the conclusion, (Grossman, 1975, Leigh, 1983, and Newhouse and Friedlander, 1980), or are unpersuasive, as in Auster, Leveson, and Sarachek (1969) which estimates regressions across states in 1960 with results that are frequently statistically insignificant and that are not robust across specifications. What is more plausible is the existence of dynamic effects, whereby mortality is positively related to transitory income and (positively) follows the business cycle, but is negatively related to permanent income. (There is a parallel here with the argument that "new" causes of mortality—cigarette smoking, obesity, lack of exercise, AIDS—first affect the rich but eventually settle down into the traditional pattern of differentially harming the poor. Income brings health risks in the short run, but ultimately the also the ability to understand and overcome them.) Even so, we must note that these results are not very robust; the correlation between the two parameter estimates is –0.80 for men and –0.78 for women. Instrumental variable results using cohort dummies as instruments (not shown here) show very different (and sometimes bizarre) patterns. Clearly, much work remains to be done, perhaps with data that are less collinear than those used here.

Table 7 also shows the results of entering education, not as average years of education, but as the proportions of the population with various levels of educational attainment ("regression 2"). These results do not differ in any major way from those in regression 1; the estimated effects of income on mortality are still positive or insignificant, and education is strongly protective, especially years of education beyond high school. Interestingly, those with "some college," shown here as 13–15 years of school, are consistently at higher risk of mortality than those with only a high-school diploma. (This effect, possibly attributable to selection, also reappears in the NLMS, albeit in a much weaker form, see Elo and Preston, 1996). Apart from this, there is no evidence here that there is a problem with using mean years of education to predict mortality. In the final regression in the table, regression 3, we repeat the first regression but with the addition of the variance of the logarithm of income per adult equivalent. The coefficients on inequality are much reduced compared with those in Table 2, typically by a factor of more than two, but the estimated protective effect remains. Allowing for the possible separate effects of education and income much reduces the size of the estimated protective effect of inequality, but it does not eliminate it.

4. Conclusions

Our original purpose was to use birth-cohort data to examine the links between mortality and inequality. Controlling for income, we find that higher inequality is associated with lower mortality, a conclusion that comes from negative association of mortality and inequality in the United States in the late 1970s and early 1980s. While it is possible that such a result has some real basis—and there are theoretical mechanisms that could produce it—it is hardly established by these results. In particular, the sign of the effect is implausible, if only because of the expected operation of Jensen's inequality, and the magnitude of the effect is quite sensitive to the way in which other variables are introduced, particularly income and education. Indeed, we suspect that the current priority should not be the investigation of the effects of inequality, but the unpacking of "socio-economic status" into its components, particularly education and income, as well as the disaggregation of mortality into its different components so as to allow them to respond to income and education in different ways. The results reported here make it clear that this is no easy task; the way in which education and income affect mortality is not the same for men and women, it is not the same for young adults as for older adults, it is different over long time periods and over the business cycle, and it is different in the cross-section from over time. We find evidence that short-term increases of income may raise the risk of mortality, particularly for young men. But in the cohort data, the longer term effects of income, or of income linked to education, are protective. Yet this evidence needs to be reconciled with the individual level data from the follow-up studies which show that, especially for men, income plays as large or larger role than education. Work on these issues has hardly begun.

5. List of works cited:-

Adler, Nancy, W. Thomas Boyce, Margaret A. Chesney, Susan Folkman, and Leonard Syme, 1993, "Socioeconomic inequalities in health: no easy solution," *Journal of the American Medical Association*, **269**, 3140–5.

Adler, Nancy E., Thomas Boyce, Margaret A. Chesney, Sheldon Cohen, Susan Folkman, and S. Leonard Syme, 1994, "Socioeconomic status and health: the challenge of the gradient", *American Psychologist*, **49**, 15–24.

Auster, Richard, Irving Leveson, and Deborah Sarachek, "The production of health, an exploratory study," *Journal of Human Resources*, **4**, 411–36.

Cohen, Sheldon, D. Tyrrell, and A. Smith, 1991, "Psychological stress and susceptibility to the common cold," *New England Journal of Medicine*, **325**, 606–12.

Cohen, Sheldon, S. Line, S. B. Manuck, B. S. Rabin, E. R. Heise, and J. R. Kaplan, 1997, "Chronic social stress, social status, and susceptibility to upper respiratory infections in nonhuman primates," *Psychosomatic Medicine*, **59**,:213–21

Deaton, Angus, 1999, "Inequalities in income and inequalities in health," Research Program in Development Studies, Princeton University, processed. (April)

Elo, Irma T., and Samuel H. Preston, 1996, "Educational differentials in mortality: United States, 1979–85," *Social Science and Medicine*, **42**, 47–57.

Fiscella, Kevin and Peter Franks, 1997, "Poverty or income inequality as predictor of mortality: longitudinal cohort study," *British Medical Journal*, **314**, 1724–8.

Fuchs, Victor R., 1974, *Who shall live? Health, economics, and social choice*. New York. Basic Books.

Fuchs, Victor R., 1989, "Comments," in John P. Bunker, Deanna S. Gombey, and Barbara Kehrer, eds., *Pathways to health: the role of social factors*, Menlo Park, Ca. Kaiser Family Foundation, pp. 226–9.

Fuchs, Victor R., 1993, "Poverty and health: asking the right questions," in David E. Rogers and Eli Ginzburg, eds., *Medical Care and the Health of the Poor*, Boulder, Co. Westview Press.

Garber, Alan M., 1989, "Pursuing the links between socioeconomic factors and health: critique, policy implications, and directions for future research," in John P. Bunker, Deanna S. Gombey, and Barbara Kehrer, eds., *Pathways to health: the role of social factors*, Menlo Park, Ca. Kaiser Family Foundation, pp. 271–315.

Grossman, Michael, 1975, "The correlation between health and schooling," in Nestor E. Terleckyj, ed., *Household production and consumtpion*

House, James S. and David R. Williams, 1995, "Psychosocial pathways linking SES and CVD," in *Report on the conference on socioeconomic status and cardiovascular health and disease*, National Heart, Lung, and Blood Institute, NIH.

Judge, Ken, 1995, "Income distribution and life expectancy: a critical appraisal," *British Medical Journal*, **311**, 1282–5

Judge, Ken, Jo-Ann Mulligan and Michaela Benzeval, 1998, "Income inequality and population health," *Social Science and Medicine*, **46**, 567–79.

Kaplan, George, Elsie R. Pamuk, J. M. Lynch, Richard D. Cohen and Jennifer L. Balfour, 1996, "Inequality in income and mortality in the United States: analysis of mortality and potential pathways," *British Medical Journal*, **312**, 999—1003.

Kennedy, Bruce P., Ichiro Kawachi, and Deborah Prothrow-Stith, 1996, "Income distribution and mortality: cross sectional ecological study of the Robin Hood index in the United States," *British Medical Journal*, **312**, 1004–7.

Lantz, Paula M., James S. House, James M. Lepkowski, David R. Williams, Richard P. Mero, and Jieming Chen, 1998, "Socioeconomic factors, health behaviors, and mortality," *Journal of the American Medical Association*, **279**, 1703–8.

Leigh, J. Paul, 1983, "Direct and indirect effects of education on health," *Social Science and Medicine*, **17**, 227–34.

Lobmayer, Peter and Richard G. Wilkinson, 1999, "Income, inequality, and mortality in 14 developed countries," Trafford Center for Medical Research, University of Sussex, processed.

Marmot, Michael G., 1994, "Social differences in health within and between populations," *Daedalus*, **123,** 197–216.

Marmot, Michael G., 1997, "Inequality, deprivation, and alcohol use," *Addiction*, **92**, S13–S20.

McIsaac, S. J., and Richard G. Wilkinson, 1997, "Income distribution and cause specific mortality," *European Journal of Public Health*, **7**, 45–53.

Newhouse, Joseph P., and Lindy J. Friedlander, 1980, "The relationship between medical resources and measures of health: some additional evidence," *Journal of Human Resources*, **15**, 200–18.

Preston, Samuel H., 1975, "The changing relation between mortality and level of economic development," *Population Studies*, **29**, 231–48.

Robert, Stephanie A., and James S. House, 1998, "Socioeconomic inequalities in health: integrating individual-, community-, and societal-level theory research," University of Michigan, processed. (forthcoming, *Handbook of Social Studies in Health and Medicine.*)

Rogot, E., P. D. Sorlie, N. J. Johnson, and C. Schmitt, eds, 1992, *A mortality study of 1.3 million persons by demographic, social, and economic factors: 1979–1985 Follow-up*, Bethesda, Md. NIH.

Ruhm, Christopher J., 1999, "Are recessions good for your health?" University of North Carolina, Greensboro, processed. (Feb.)

Sapolsky, Robert M., 1993, "Endocrinology alfresco: psychoendocrine studies of wild baboons," *Recent Progress in Hormone Research*, **48**, 437–68.

Sapolsky, Robert M., 1994, *Why zebras don't get ulcers*, New York. Freeman.

Wilkinson, Richard G., 1996, *Unhealthy societies: the afflictions of inequality*, London, Routledge.

Wilkinson, Richard G., 1997, "Commentary: Income inequality summarises the health burden of individual relative deprivation," *British Medical Journal*, **314**, 1727–8.

Wilmoth, John, 1999, *The Berkeley Mortality Database*, accessible at http://demog.berkeley.edu/wilmoth/mortality

Appendix

In the text, we make repeated use of a standard result from the normal distribution, which is stated here for convenience. In words, if two variable are jointly normally distributed, the conditional expectation of one given the other—the regression function—is linear and homoskedastic, with coefficients and residual variance equal to the coefficients and residual variance of a large sample OLS regression. Formally, suppose that

$$\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} \sim N\begin{pmatrix} \mu_1 & \sigma_{11} & \sigma_{12} \\ \mu_2 ' & \sigma_{21} & \sigma_{22} \end{pmatrix} \tag{A.1}$$

then

$$E(x_1 | x_2) = a + bx_2 \tag{A.2}$$

$$a = \mu_1 - \frac{\sigma_{12}}{\sigma_{22}} \mu_2 \tag{A.3}$$

$$b = \frac{\sigma_{12}}{\sigma_{22}} \tag{A.4}$$

and for the variance,

$$V(x_1 | x_2) = \sigma_{11} - \frac{\sigma_{12}^2}{\sigma_{22}}. \tag{A.5}$$

The same results hold for the expectation of x_2 conditional on x_1 with “1” and “2” transposed.

Table 1: Log odds of dying as a function of income, by age group. Coefficients on mean ln(income/adult equivalent), t-statistics in parentheses

	Men						Women					
	OLS		IV years of education		IV years of education cohort dummies		OLS		IV years of education		IV years of education cohort dummies	
No time trend included												
All age groups	–0.559	(24.1)	–0.923	(29.1)	–0.826	(28.6)	–0.528	(27.4)	–0.840	(31.1)	–0.838	(32.7)
ages 25–39	0.452	(5.1)	–0.940	(3.2)	0.646	(4.7)	–0.362	(6.1)	–0.995	(9.7)	–0.850	(9.8)
ages 40–54	–0.770	(17.8)	–1.046	(19.6)	–0.996	(19.6)	–0.850	(23.5)	–1.114	(24.9)	–1.081	(25.3)
ages 55–69	–0.941	(19.4)	–1.413	(21.1)	–1.311	(21.2)	–0.322	(10.8)	–0.628	(14.4)	–0.497	(13.4)
ages 70–85	–0.430	(19.0)	–0.637	(20.2)	–0.603	(20.9)	–0.370	(14.6)	–0.671	(16.8)	–0.637	(17.1)
Time trend included												
All age groups	–0.281	(8.8)	–1.224	(15.6)	–0.877	(15.1)	–0.125	(5.4)	–0.619	(9.8)	–0.686	(13.8)
ages 25–39	0.168	(1.8)	–1.363	(4.1)	–0.189	(1.0)	–0.174	(2.7)	–1.150	(7.2)	–0.809	(7.1)
ages 40–54	–0.261	(3.9)	–1.379	(6.7)	–0.918	(7.0)	–0.022	(0.5)	–0.547	(3.6)	–0.466	(4.9)
ages 55–69	0.107	(5.9)	0.101	(2.4)	0.104	(3.2)	0.207	(9.3)	0.275	(4.5)	0.380	(8.9)
ages 70–85	0.035	(1.7)	–0.118	(1.4)	0.037	(0.8)	0.059	(2.8)	–0.194	(1.8)	–0.035	(0.6)

Notes: Each regression includes the mean of the logarithm of household income per adult equivalent (mean ln(y/ae)), and a set of age dummies. The regressions are estimated for the full sample, and for subsets of cohorts in different age groups. Each cell in the table reports a coefficient on the mean of log of income per adult equivalent from a single regression. The IV estimates instrument mean ln(y/ae) with the mean years of education (in the columns marked "educ") or with mean years of education and a set of birth-cohort dummies (in the columns marked "years of education, cohort dummies")

Table 2: Log odds of dying as a function of income and income inequality, by age group.

	Men				Women			
	OLS		IV		OLS		IV	
Inequality measure is variance of ln(*y*/*ae*)								
All age groups								
mean ln(*y*/*ae*)	–0.520	(22.7)	–0.606	(8.1)	–0.479	(26.2)	–0.417	(4.5)
var ln(*y*/*ae*)	–0.157	(8.7)	–0.655	(4.9)	–0.180	(13.6)	–0.656	(5.1)
Ages 35-59								
mean ln(*y*/*ae*)	–0.621	(14.6)	–1.368	(4.9)	–0.726	(26.2)	–0.307	(1.5)
var ln(*y*/*ae*)	–0.236	(8.0)	0.353	(1.0)	–0.240	(11.4)	–1.066	(3.9)
Ages 60-85								
mean ln(*y*/*ae*)	–0.506	(21.6)	–0.594	(9.4)	–0.298	(15.9)	–0.380	(5.7)
var ln(*y*/*ae*)	–0.123	(5.7)	–0.475	(3.3)	–0.084	(5.7)	–0.375	(3.0)
Inequality measure is gini coefficient of *y*/*ae*								
All age groups								
mean ln(*y*/*ae*)	–0.505	(20.3)	–0.596	(8.0)	–0.440	(21.7)	–0.631	(12.4)
gini *y*/*ae*	–0.525	(5.5)	–2.910	(5.1)	–0.819	(10.5)	–1.522	(5.2)
Ages 35-59								
mean ln(*y*/*ae*)	–0.504	(10.6)	–1.370	(6.6)	–0.625	(18.9)	–0.588	(3.4)
gini *y*/*ae*	–1.563	(8.3)	1.476	(1.4)	–1.398	(9.8)	–2.624	(2.8)
Ages 60-85								
mean ln(*y*/*ae*)	–0.518	(20.3)	–0.694	(15.3)	–0.308	(15.1)	–0.435	(7.0)
gini *y*/*ae*	–0.154	(1.5)	–0.870	(2.4)	–0.105	(1.2)	–1.514	(2.2)

Notes: Each regression includes the mean of the logarithm of household income per adult equivalent (mean ln(*y*/*ae*)), a set of age dummies, and an inequality measure. The IV estimates instrument mean ln(*y*/*ae*) and the inequality measure with the mean years of schooling, and the fraction of people in education categories. The education categories are years of schooling equal to 5-8, 9-11, 12, 13-15, and 16 or more. Years of education from 0 to 4 is the omitted category.

Table 3: Poverty; interactions between income and inequality

	Men				Women			
	OLS		IV		OLS		IV	
mean ln(y/ae)	–0.646	(25.0)	–1.101	(24.3)	–0.670	(33.6)	–0.970	(34.0)
fraction poor	–0.825	(7.1)	–1.774	(5.4)	–1.181	(15.2)	–1.447	(7.0)
mean ln(y/ae)	–0.580	(21.1)	–0.688	(2.9)	–0.603	(28.7)	–0.700	(4.9)
var ln(y/ae)	–0.125	(6.4)	–0.552	(1.8)	–0.117	(8.4)	–0.333	(1.9)
fraction poor	–0.490	(3.9)	–0.321	(0.4)	–0.884	(10.6)	–0.853	(2.3)
mean ln(y/ae)	–0.263	(5.5)	–0.274	(1.8)	–0.215	(5.2)	–0.261	(2.4)
var ln(y/ae)	4.165	(5.9)	8.191	(2.3)	3.769	(6.7)	3.215	(1.7)
mean*var	–0.474	(6.1)	–0.929	(6.1)	–0.438	(7.0)	–0.413	(7.0)
∂/∂mean ln(y/ae)	–0.539	(23.6)	–0.815	(7.7)	–0.485	(27.0)	–0.515	(5.6)
∂/∂var ln(y/ae)	–0.160	(9.0)	-0.299	(1.6)	–0.188	(14.3)	–0.518	(4.0)

Notes: Each regression also includes a set of age dummies. The IV estimates instrument mean ln(y/ae), poverty and inequality with the mean years of schooling, and the fraction of people in different education categories. The education categories are years of schooling equal to 5-8, 9-11, 12, 13-15, and 16 or more. Years of education from 0 to 4 is the omitted category. The derivatives are evaluated at the sample means of var ln(y/ae) and mean ln(y/ae), respectively. The means for var ln(y/ae) are 0.583 (mean) and 0.617 (women). The means for mean ln(y/ae) are 9.135 (men) and 9.039 (women).

Table 4: The effects of education and income on mortality in the NLMS

	Men		Women	
ln(y/ae)	–0.3524	(10.8)	–0.2620	(6.3)
years of education	–0.0374	(6.2)	–0.0483	(6.0)
ln(y/ae)	–0.3254	(9.1)	–0.1968	(4.3)
years of education	–0.0121	(1.8)	–0.0336	(3.8)

Notes: Coefficients from logit regressions in which dependent variable is whether or not the respondent died within 365 days of interview. A full set of age dummies, one for each year from 25 to 85, are included but not shown. There are 161,472 males in the sample and 183,282 females. The income and education variables are entered separately (the first two rows) and then together (last two rows.)

Table 5: Effects of income and schooling on male mortality, NLMS (Logits)

	Died in 0–1 years		Died in 1–2 years		Died in 2–5 years		Died in 5–9 years	
			Men aged 25–59 at time of survey					
ln(y/ae)	–0.456	(8.14)	–0.397	(7.32)	–0.415	(14.01)	–0.325	(13.43)
year school	–0.072	(5.70)	–0.054	(4.47)	–0.076	(11.62)	–0.063	(11.85)
ln(y/ae)	–0.390	(6.25)	–0.357	(5.94)	–0.328	(9.95)	–0.248	(9.24)
years school	–0.035	(2.50)	–0.021	(1.55)	–0.045	(6.27)	–0.040	(6.81)
observations	123,806		123,298		122,740		120,811	
number of deaths	508		558		1,929		3,059	
			Men aged 60–85 at time of survey					
ln(y/ae)	–0.302	(7.62)	–0.197	(5.10)	–0.224	(9.43)	–0.202	(9.29)
years school	–0.028	(4.06)	–0.034	(5.14)	–0.028	(6.73)	–0.026	(6.62)
ln(y/ae)	–0.285	(6.54)	–0.138	(3.23)	–0.191	(7.26)	–0.172	(7.11)
years school	–0.007	(0.88)	–0.024	(3.26)	–0.013	(2.91)	–0.012	(2.75)
observations	37,666		36,123		34,514		29,705	
number of deaths	1,543		1,609		4,809		6,061	

Notes: For men aged 25–59, age at the year of the survey was also included. For men aged 60–85, a complete set of age dummies were included. The sample for each logit includes individuals who either died in the time period at the head of column or died later; individuals who died earlier than the time period specified are excluded.

Table 6: Effects of income and schooling on female mortality, NLMS (Logits)

	Died in 0–1 years		Died in 1–2 years		Died in 2–5 years		Died in 5–9 years	
				Women aged 25–59 at time of survey				
ln(y/ae)	–0.409	(5.75)	–0.340	(4.78)	–0.344	(9.03)	–0.332	(11.41)
years school	–0.076	(4.13)	–0.060	(3.26)	–0.066	(6.74)	–0.065	(8.66)
ln(y/ae)	–0.348	(4.41)	–0.296	(3.75)	–0.287	(6.80)	–0.276	(8.54)
years school	–0.037	(1.82)	–0.027	(1.34)	–0.035	(3.16)	–0.035	(4.17)
observations	134,355		134,041		133,721		132,581	
number of deaths	314		320		1,140		1,995	
				Women aged 60–85 at time of survey				
ln(y/ae)	–0.190	(3.75)	–0.074	(1.52)	–0.174	(6.50)	–0.133	(6.01)
years school	–0.042	(4.67)	–0.017	(1.96)	–0.024	(4.95)	–0.019	(4.56)
ln(y/ae)	–0.118	(2.16)	–0.044	(0.84)	–0.145	(5.01)	–0.111	(4.60)
years school	–0.034	(3.50)	–0.014	(1.49)	–0.014	(2.64)	–0.011	(2.37)
observations	48,927		47,900		46,789		42,789	
number of deaths	1,027		1,111		4,000		6,202	

Notes: For women aged 25–59, age at the year of the survey was also included. For women aged 60–85, a complete set of age dummies were included. The sample for each logit includes individuals who either died in the time period at the head of column or died later; individuals who died earlier than the time period specified are excluded.

Table 7: Mortality, income and education: cohort data

	Men						Women					
	All ages		Ages 35-59		Ages 60-85		All ages		Ages 35-59		Ages 60-85	
Regression 1:												
mean ln(*y*/*ae*)	0.088	(2.9)	0.285	(4.7)	0.002	(0.1)	–0.022	(0.9)	–0.001	(0.0)	0.090	(4.71)
mean yrs education	–0.117	(26.2)	–0.192	(19.6)	–0.088	(22.7)	–0.097	(25.6)	–0.159	(23.2)	–0.071	(27.7)
Regression 2:												
mean ln(*y*/*ae*)	0.079	(2.8)	0.173	(3.2)	–0.018	(0.7)	0.065	(2.9)	0.019	(0.5)	0.081	(4.2)
5-8 years school	0.205	(1.7)	–0.236	(0.5)	0.330	(3.8)	–0.356	(3.1)	–0.300	(0.8)	–0.402	(5.2)
9-11 years school	0.142	(1.3)	–1.430	(3.4)	0.245	(2.9)	–0.003	(0.0)	–0.319	(0.9)	–0.661	(9.5)
12 years school	–0.463	(4.5)	–2.029	(5.1)	–0.160	(2.1)	–0.704	(7.7)	–1.452	(4.5)	–0.749	(12.2)
13-15 years school	–0.107	(1.0)	–0.996	(2.5)	–0.483	(5.7)	–0.513	(5.2)	–0.779	(2.4)	–0.798	(10.4)
16+ years school	–1.502	(12.8)	–2.911	(7.4)	–0.795	(8.4)	–1.782	(17.1)	–2.255	(6.8)	–1.045	(12.7)
Regression 3:												
mean ln(*y*/*ae*)	0.077	(2.5)	0.267	(4.3)	–0.003	(0.1)	–0.053	(2.1)	–0.043	(1.0)	0.087	(4.6)
mean yrs education	–0.112	(24.4)	–0.186	(16.9)	–0.086	(21.7)	–0.087	(21.6)	–0.148	(18.3)	–0.070	(26.4)
var ln(*y*/*ae*)	–0.061	(4.0)	–0.027	(1.0)	–0.051	(3.2)	–0.079	(6.5)	–0.052	(2.7)	–0.011	(1.1)

Notes: Each regression includes the mean of the logarithm of household income per adult equivalent (mean ln(*y*/*ae*)), and a set of age dummies. The regressions are estimated for the full sample, and for subsets of cohorts in different age groups.

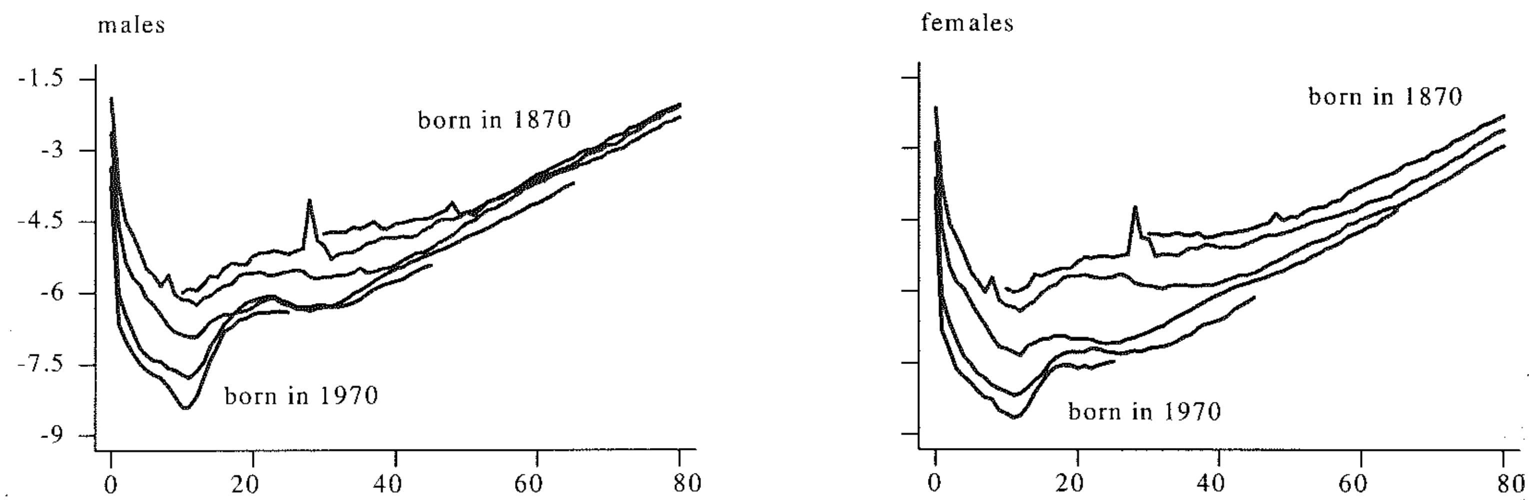

Figure 1: Log odds of mortality, every 20th cohort

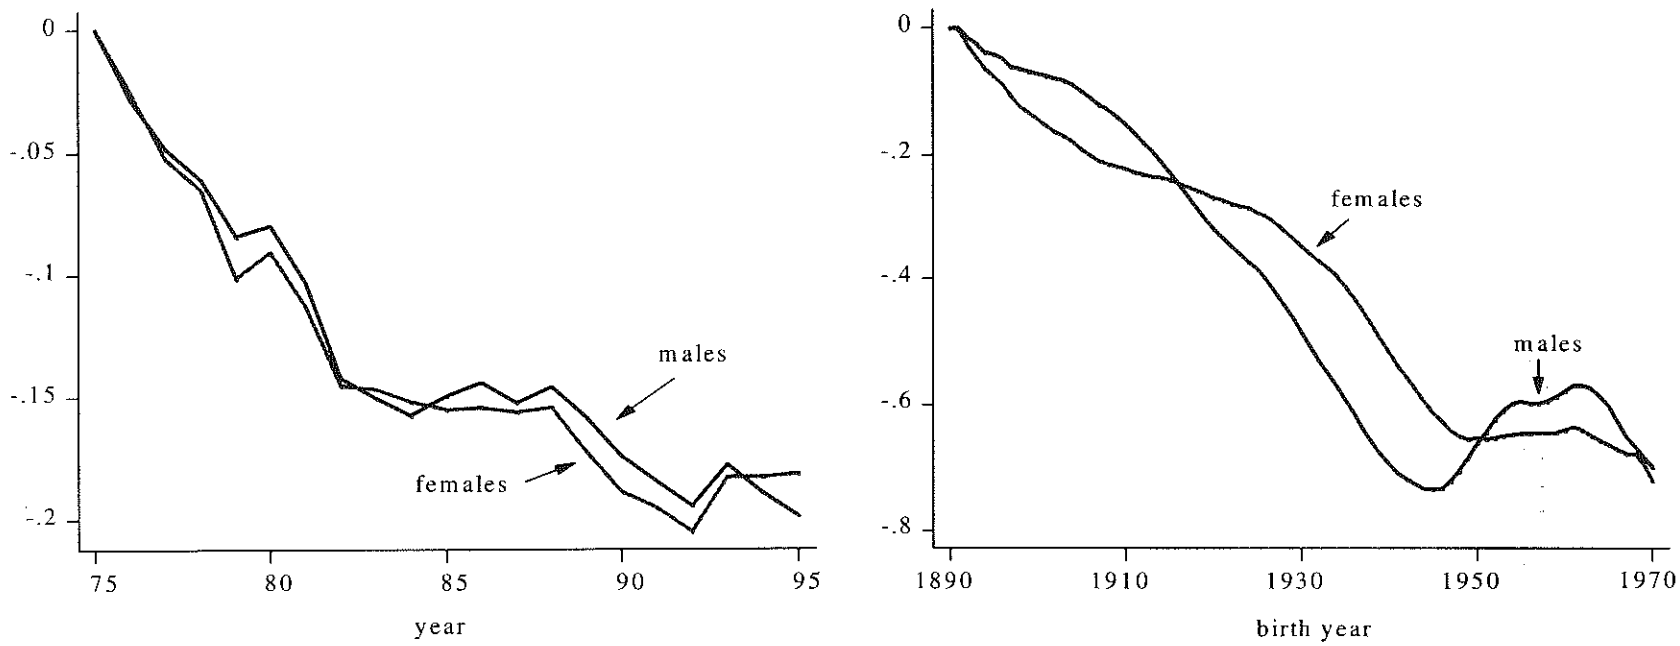

Figure 2: Year and cohort effects in log odds of mortality, from regressions of log odds on age and year effects, and on age and cohort effects.

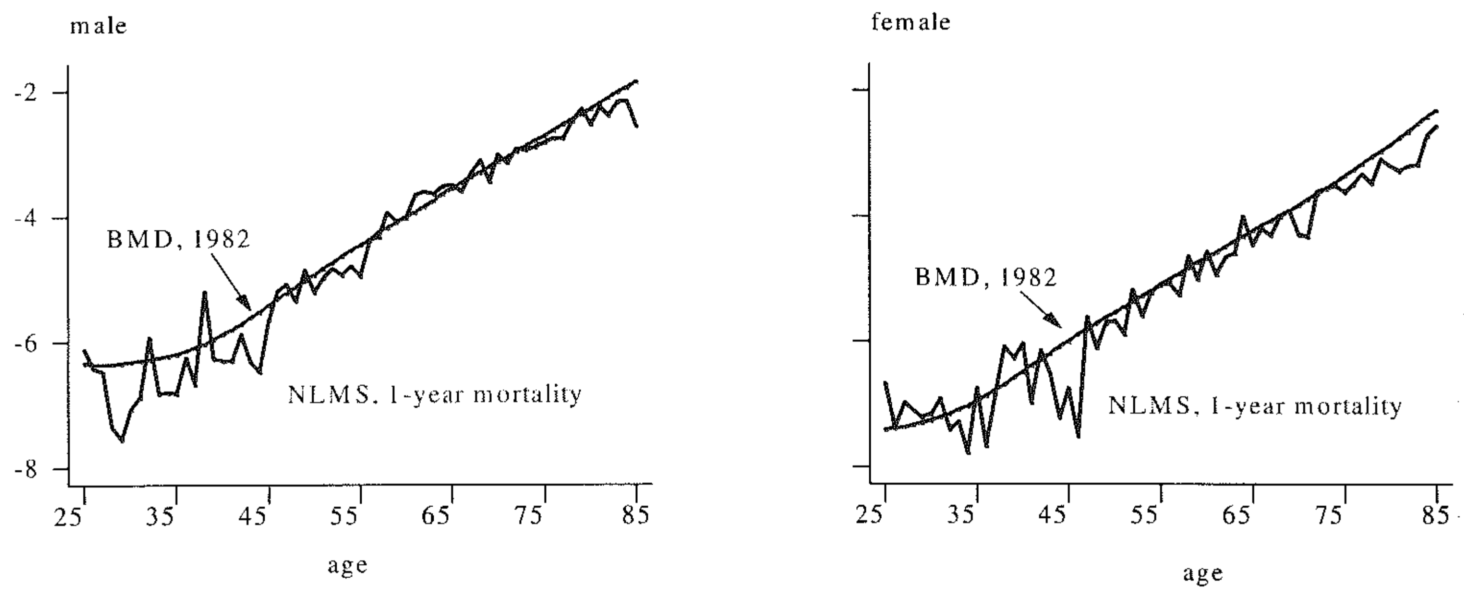

Figure 3: Log odds of mortality, NLMS and BMD

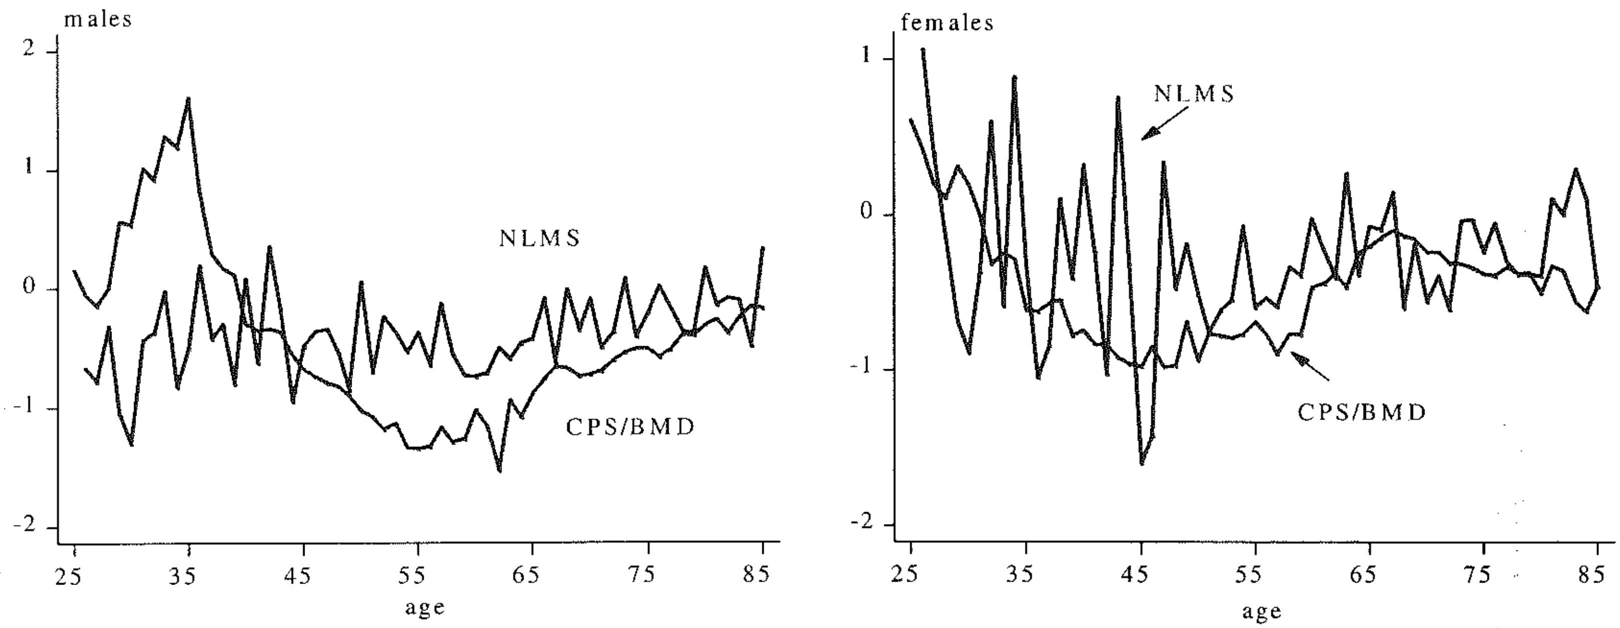

Figure 4: Effects of income on mortality at each age, NLMS and CPS/BMD

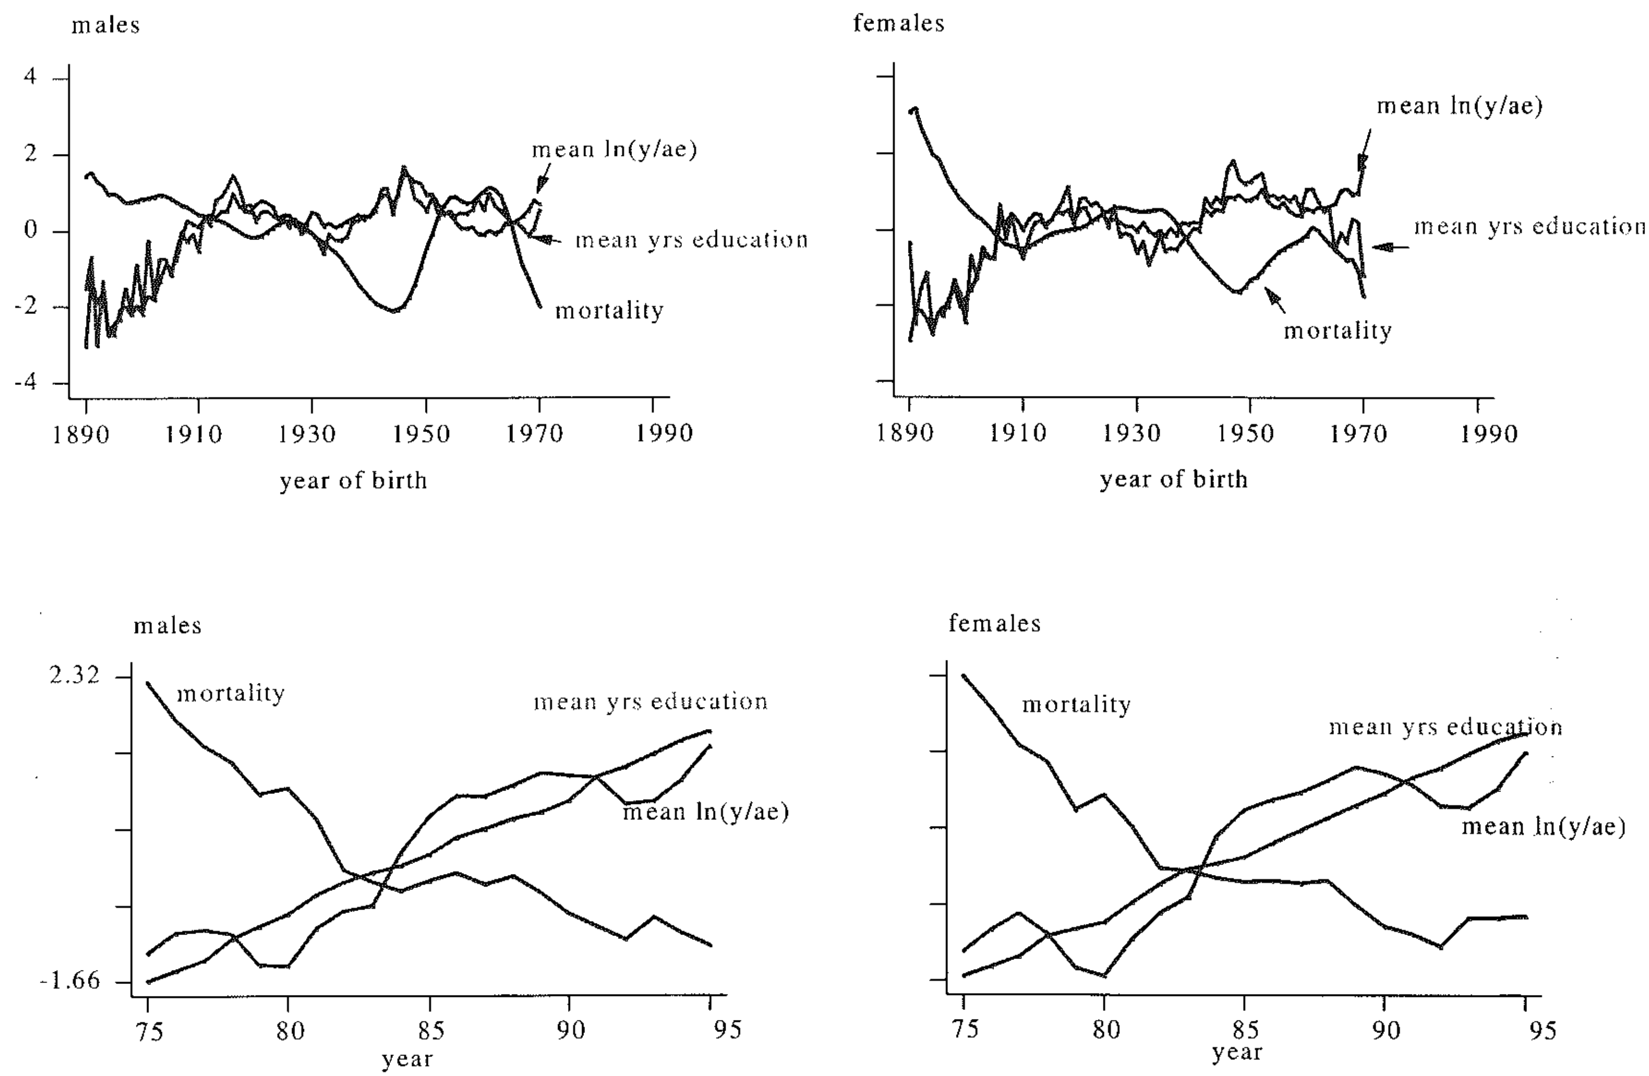

Figure 5: Cohort and year averaged residuals from regressions on age dummies (variables standardized to have mean=0, std.dev.=1)

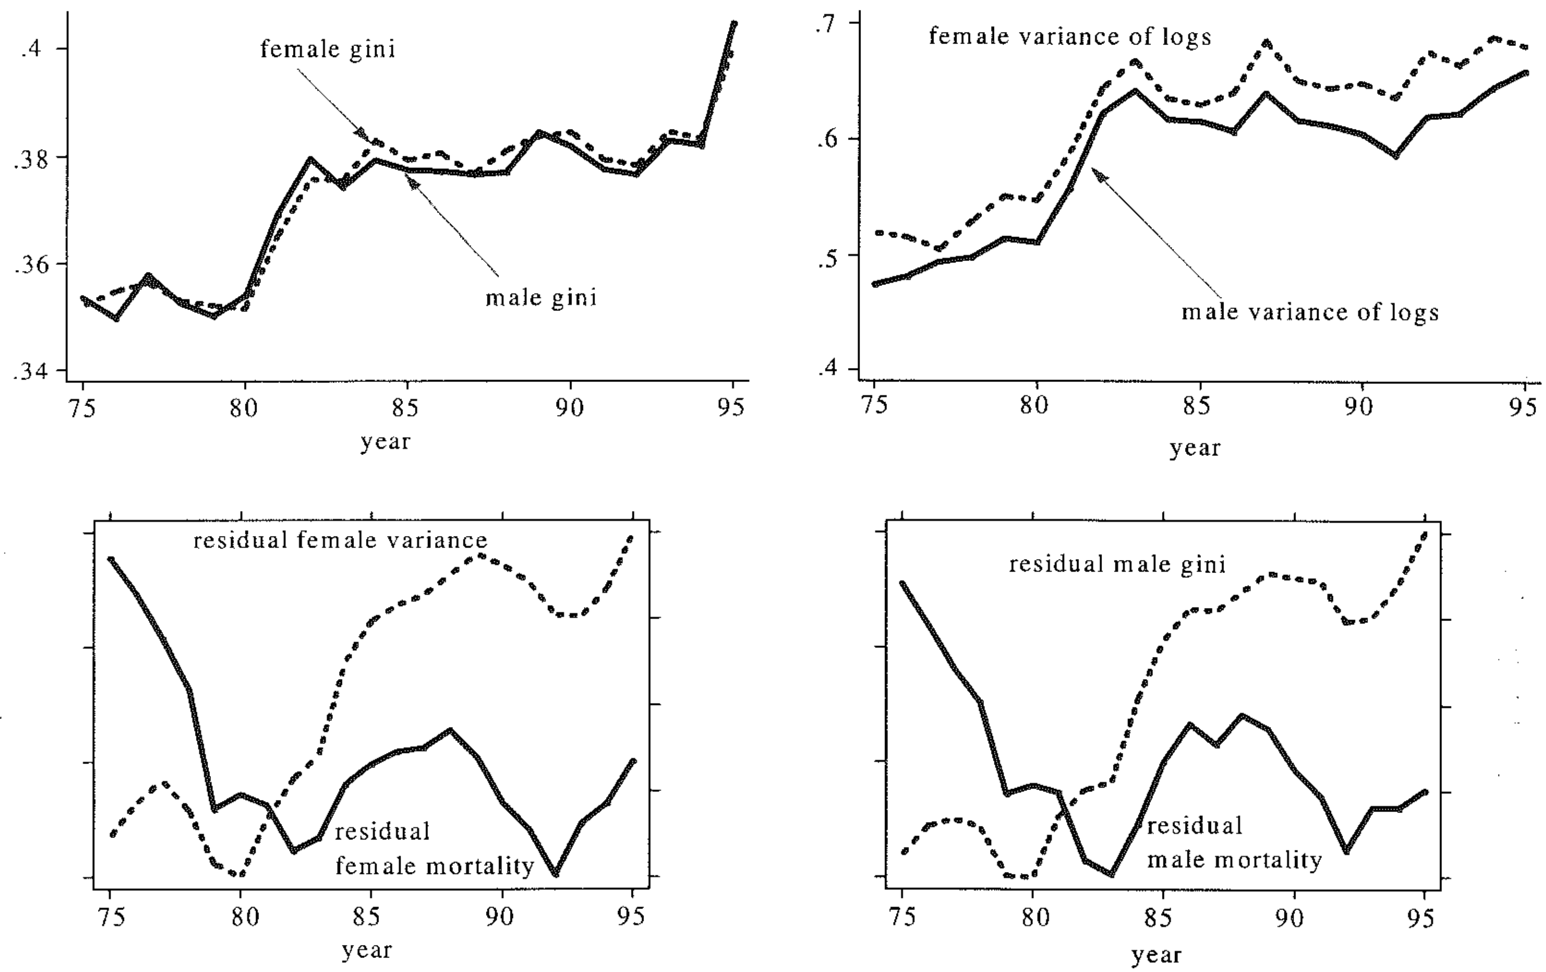

Figure 6: Inequality and mortality