

Inequalities in income and inequalities in health

Angus Deaton

Research Program in Development Studies
Princeton University

This version
May 19th, 1999

Prepared for the conference on “Increasing inequality in America,” The Bush School of Public Policy, Texas A. and M. University, March 12–13, 1999. Research for this paper is supported by the John D. and Catherine T. MacArthur Foundation’s Network on Inequality and Poverty in Broader Perspectives. I should like to thank the MacArthur Foundation and the National Institute of Aging for Financial Support, and Anne Case, Noreen Goldman, Alan Krueger, Christina Paxson, Burt Singer, James Smith, and James Trussell for comments.

ABSTRACT

What is inequality in health? Are economists' standard tools for measuring income inequality relevant or useful for measuring it? Does income protect health and does income inequality pose a hazard to health? In this paper, I discuss two different concepts of health inequality and relate each of them to the literature on inequality in income. I propose a model of reference groups, in which each individual's health is related to his or her status within the group as measured by income relative to the group mean. In such a model, increases in income inequality, whether within groups, or between them, has no effect on average health. Even so, the slope of the relationship between health and measured income, the "gradient," depends on the ratio of between to within-group income inequality. The model is extended to allow income inequality to play a direct positive or negative role in determining health status. Empirical evidence on cross-country income inequality and life expectancy within the OECD, and on time-series evidence for the U.S., Britain, and Japan, provide little support for the idea that inequality is a health hazard at the national level. Data on birth cohorts for the US observed from 1981 to 1993 show no relationship between mortality and income inequality. However, there is a well defined health gradient in these data, and its slope increases with increases in each cohort's income inequality.

0. Introduction.

In much of applied welfare economics, income—real income, consumption, or consumption per equivalent—serves as our measure of individual welfare. But the goods and services provided by income are not all that there is to well-being. Health is not only instrumental in enabling people to earn a living, and to enjoy the fruits of their labors, but is an important element of well-being in its own right. Health status is correlated with income, both for individuals within nations, and across nations in aggregate. But the correlation is far from perfect so that looking at health leads to different assessments of well-being than come from looking only at income. Much the same can be said for education, and such considerations have led to simple measures such as the UN's Human Development Index, a composite of life-expectancy, income, and literacy. In this paper, I consider only two of the three components, health and income. In particular, I am concerned with what it means to talk about inequality in health, and whether, according to some useful definition of the concept, health inequality in the United States is rising in tandem with the rise in income inequality. I also investigate the possibility that income inequality itself is a health hazard, a hypothesis advocated by Richard Wilkinson (1996).

The first section of the paper is concerned with the concept of health inequality, with alternative definitions, and with the evidence that health inequalities in the United States are increasing. The second section turns to the Wilkinson hypothesis. In an attempt to develop its theoretical underpinnings, I present two simple models in which there is a link between aggregate inequality and individual health, in only one of which is inequality a fundamental risk factor in its own right. It is widely understood that a nonlinear (typically concave) relationship between health and income at the individual level will generate an aggregate relationship in which average

health depends (negatively) on the degree of inequality. But I also show that if health status depends not on income, but on income relative to the incomes of members of some reference group, then the relationship between income and health, unconditional on information about the reference group—and such information is typically unavailable—depends on the relative size of within- and between-group inequality. The model can readily be extended to incorporate a more fundamental role for inequality in the determination of health status, for example by making each individual's health within the reference group depend on the shortfall of individual income from that of the richest person in the group. But I argue that such mechanisms can just as well generate a positive as a negative link between inequality and health. That individual health depends on relative income is consistent has no direct implication for the relationship between inequality and average health. Section 3 presents some empirical evidence on mortality in the United States at a hitherto unexplored level of aggregation, that of a birth cohort. By matching cohort mortality rates to cohort income and income inequality data from the Current Population Surveys, I can test whether the rising income inequality in the 1980s played any role in the contemporaneous slowdown in mortality decline. Although the inequality hypothesis is hardly well-enough developed to permit a sharp and convincing test, I find little or no evidence of a direct link between income inequality and mortality. However, the cohort data suggest that increasing inequality increases the slope of the regression of mortality on income, a finding that is predicted by the theory if the increase in income inequality acts to increase inequality within reference groups by more than inequality between them. Section 4 concludes and discusses a number of important caveats.

1. Rising health inequality?

A good way to approach health inequality is to start with income inequality, and to ask whether the theoretical and measurement structure of the latter can be transferred to the former. Measures of income inequality are measures of dispersion of the (univariate) distribution of income across persons. Questions of why such quantities are of interest, or whether some are of more interest than others, can be answered through the theoretical apparatus developed by Anthony Atkinson (1970) and Amartya Sen (1973). Inequality aversion, or a preference for a more equal distribution is coded into a social welfare function according to which mean-preserving but equalizing transfers increase social welfare or, alternatively, one in which there is diminishing (social) marginal utility to income. As Atkinson showed, these formulations lead to an aggregate measure of welfare which can be thought of as the product of mean income and income *equality*, which is the complement of inequality, and is a number between zero and one.

Health promotes well-being, just as income promotes well-being, and some people have better health than others; statements like A is healthier than B are meaningful in much the same way as are statements such as A has more income than B. But an immediate problem is measurement; we have a cardinal measure of income, but no comparable measure of health status. For populations, life-expectancy is a useful statistic with convenient properties, but is much less useful at the individual level. Life expectancy at a given age in a given year is not the expected years of life of a person of that age in that year; period estimates of life expectancy are based on age-specific mortality rates at a fixed moment in time, rates that can be expected to change (typically decline) by the time the individual reaches those ages. The expected lifespan for each individual would be a good tool for measuring well-being, but its calculation among the living requires an

assessment of health status with a view to assessing individual mortality probabilities in the future. The measurement of health status is extremely difficult, except among the elderly among whom functional impairments are sufficiently widespread and apparent to allow useful calibration. Self-reported measures of self-reported health status (my health is “poor,” “fair,” “good,” “very good” or “excellent”) have been shown to predict later health conditions, but estimates of inequality based on such reports are typically not invariant to the wide range of equally reasonable ways of scoring or scaling the responses. Discussions of how to measure inequality using qualitative data can be found in Deaton and Paxson (1998b) and Allison and Foster (1998).

Suppose that there were no measurement problems, and that for each individual we could readily calculate a measure of life expectancy. One could imagine doing this long after everyone has died, so that we have an actual distribution of years lived for a birth cohort, or we might be prepared to accept the period life-tables and calculate for each person alive now an expected age at death. Inequality in years lived could then be calculated, as for 32 developed countries in Julian Le Grand (1987). But as Le Grand recognizes, the axioms that underlie inequality measurement are not obviously applicable in this context. For example, it is unclear how many people would assent to the proposition that society would be better-off if a fifty year-old died at 49 instead of 50 in order to extend the life of a 45 year-old by one year. Unlike income, where policies for effecting transfers can readily be imagined, transfers of life are not readily linked to health policy, so that it is unclear that reduction in inequality in health outcomes is a worthy (or feasible) target for policy.

There is another literature on “health inequalities” which defines inequality quite differently. In the public health, psychology, sociology, demography, and epidemiology literatures, inequali-

ties in health are taken to refer to the differences in health across different socio-economic groups, typically defined by income, occupation, or education. That mortality was higher for members of lower status occupations was noted in Britain from the middle of the last century, and the failure of this “gradient” to vanish in response to the introduction of the National Health Service after the Second World War has been the impetus for an enormous amount of subsequent research as well as political acrimony, see Sally McIntyre (1997) for a fine review.

Although much of the best work on the gradient is still done in Britain, particularly the Whitehall study run by Michael Marmot and his collaborators, socioeconomic differences in health status are also well-established in the United States. For example, the National Longitudinal Mortality Study has merged subsequent death certificates back into Current Population and Census data from the late 1970s until the mid-1980s, and has generated very large samples for examining differences in mortality by income. Using these data, Rogot et al (1992) calculate life expectancy by age for seven family income groups. Comparing the bottom income group, defined as those with less than \$5,000 of family income in 1980, with the top group, which had more than \$50,000, life expectancy for men at age 25 was 43.6 years at the bottom as opposed to 53.6 years at the top. Although the absolute difference is smaller for older men, the proportional difference remains more or less constant, 26.2 versus 39.0 years at age 45, and 13.3 versus 17.2 years at age 65. Family income differences are less important for women’s mortality than for men’s mortality, but the gradient is still apparent for women. Life-expectancy appears to increase monotonically with income; excess mortality is not associated only with poverty. Nor can the gradient readily be removed by controlling for other factors, such as race or smoking behavior. Causality almost certainly runs both ways, from health to income—for example through the

effects of ill-health on employment—but there is fairly wide agreement that the effects that run from income to health are a major part of the story. And while income almost certainly is standing in for other factors such as education, both factors are separately important, Elo and Preston (1996).

Most commentators see these health inequalities as deeply offensive, more so than the economic and social inequalities to which they are related. Their elimination is seen as an urgent priority for public health policy, and some governments and international agencies have accepted the reduction of health inequalities as target. For example, the countries of the European Region of the World Health Organization unanimously adopted a resolution in 1984 that “by the year 2000 the actual differences in health status between countries and between groups within countries should be reduced by at least 25 percent.” Many people who are prepared to accept inequality in the allocation of goods as a (possibly) necessary evil are not prepared to accept similar inequalities in health outcomes. However, since health inequalities appear to be much the same whether or not health care is provided through the market, for example in the U.S. versus Britain, and to be much the same whether diseases are treatable or not, the remedy is not the provision of access to health care without reference to financial resources, Adler et al (1993). Indeed, the literature makes a good case that inequalities in mortality seem to have relatively little to do with better access to health care by better-off people. Health status may not be much affected by health care, and health inequalities may be as deeply rooted and as intractable as the social inequalities within which they are set.

One issue in the current debate is the claim that health inequality is increasing, not in the sense that the univariate distribution of health status is widening, but in the sense that the gradi-

ent is steepening, so that the same difference in socio-economic status, income or education, is now associated with a larger difference in the probability of death. (Univariate measures of health inequality in terms of years lived have decreased with reductions in infant mortality; because the univariate distribution is bimodal, reductions in the mass at the lower mode are a powerful force in reducing inequality.) Studies by Feldman et al (1989) and by Pappas et al (1993) have concluded that the differences in mortality across educational groups are larger than those in 1960 documented in the pioneering study of Kitagawa and Hauser (1973). Deaton and Paxson (1998a) find positive time trends in the variance of self-reported health status and in its correlation with income, so that the gradient between (this measure of) health status and income has been trending upward, something that is true for both men and women. The best estimates of the changes in the relationship between mortality and education come from Preston and Elo (1995), again using the NLMS. They confirm that for white males, moving from the bottom to the top of the education distribution reduced the standardized death rate in 1979–85 by three times as much as it did in 1960. They find no similar effect for females. Among the possible causes for the changes, the authors point out that the changing distribution of heart disease over education groups accounts for much of the change, and argue that improvements in prevention and treatment may have diffused more rapidly among more educated men. Other potential causes are the widening income difference associated with educational differences—the increasing rate of return to education—and the decrease in income differences between men and women.

Both of these arguments point to a protective role for income that operates independently of education, and implicate in health inequalities the same factors (particularly the increase in the returns to education) that are usually identified as contributing to increases in income inequality.

Such a mechanism provides a link between income inequality and both concepts of health inequality; growing income differences are associated both with an increase in spread in the univariate distribution of health, as well as with an increase in the gradient linking education and health. In the simplest model, we might write an earnings function as

$$y_i - \bar{y} = \beta (S_i - \bar{S}) + \varepsilon_i \quad (1)$$

where y is the logarithm of income, and S years of schooling, coupled with a health status equation

$$H_i - \bar{H} = \alpha (y_i - \bar{y}) + \eta_i \quad (2)$$

where H is some measure of health status that is (negatively) linked to mortality. As the rate of return to education β increases with a fixed distribution of schooling, the distribution of income widens, as does the distribution of health. The slope of the health to education gradient is $\alpha\beta$, which also increases with increases in the rate of return to schooling, β . This model works much better for men than for women, for whom β has also risen. Other factors at work include the fact that α appears to be smaller for women than men, and that the feedback from health to income is more important for men whose participation in the labor force is higher.

Health inequality, like income inequality, is not always undesirable. In the Atkinson framework, social welfare is the product of mean income and equality so that even if rises in inequality pull down the second factor, the effect may be offset by increases in the first. We must judge the outcome by both factors, not just one. Focusing on an increase in inequality by itself is no less (and no more) incorrect than focusing on mean income and ignoring distribution. The major concern with income inequality in the US today is the lack of growth of real wages for all but the top of the distribution, not the increase in earnings inequality per se. A pertinent example is provided

by the consequences of an increase in the rate of return to education, whether generated by skill-biased technical change or, to take a more hypothetical example, by an increase in the quality of schools. Such a change raises wages for everyone and encourages people to become better educated. At the same time, because some people get more schooling than others, it raises inequality.

Similar considerations appear to apply to much of the change in health inequality. If new knowledge about preventive methods for reducing heart disease—particularly exercise and diet—are adopted more readily by people with higher incomes (and education), the new knowledge has brought about a Pareto improvement, in spite of the increase in health inequality. Another example is infant mortality among whites in the U.S., which is roughly twice as high among the least educated as it is among the most educated mothers, a difference that is almost entirely explained by differential rates of cigarette smoking, even though both sets of mothers understand the risks, see Ellen Meara (1998). Once again the knowledge that became widely available and widely credible after the publication of the Surgeon General’s report on smoking in 1964 has led to a (Pareto) improvement in public health, but one that has been unequally distributed. Of course, the reduction of infant mortality rates among poorly educated women remains a major priority for public health, though improvements, like improvements in mortality from heart disease, are likely to require, not technical improvements in medicine or even better access to health care, but some understanding of why low incomes and poor education make it so difficult to benefit from the relevant knowledge. But it is the lack of improvement in health among the poorly educated—like the lack of growth in their real wages—that is the central policy issue, not the increase in inequality per se. We would surely not condemn the Surgeon General’s report on the grounds that it led to an increase in health inequality.

2. Inequality in incomes and the level of health

A different link between incomes and health has been proposed by Wilkinson (1996) who sees income inequality itself as a principal determinant of poor health. This is quite different from the linkage discussed in the previous paragraph where the *level* of income was tied to the *level* of health, so that income *inequality* was tied to health *inequality*. The Wilkinson hypothesis is that inequality itself poses a risk to health that, at any given age, those who live in a more unequal society have a higher probability of death. Although Wilkinson does not rule out the possibility that average health is affected by income inequality through the aggregation of a concave relationship between risk and mortality, he makes it clear that the hypothesis goes further and asserts that inequality itself is a risk to each individual. There are some good theoretical reasons why such a relationship might exist; for example, the proximity of rich and poor might lead to crime and violence, as in many South African cities today. More generally, the precise mechanisms are hard to pin down, and in some of the discussions in the literature, inequality seems to be standing proxy for a wide range of social ills; indeed, in the non-economics literature, the terms “inequality” and “poverty” are not always carefully delineated. Some mechanisms are hard to understand. How, for example, would an increase in property taxes for wealthy homeowners in Phoenix, Arizona lead to a decrease in the probability of a middle-aged man in New York dying from heart disease?

Yet the potential importance of establishing a link between mortality and inequality makes it worth serious theoretical and empirical investigation. For example, a more progressive income tax could be defended on public health grounds, or as argued by Lynch, Kaplan and Salonen (1997, p. 809), “efforts to reduce socioeconomic inequalities in health must recognize that

economic policy is public health policy.” The creation of individual saving accounts as part of social security reform would cause greater inequality at the time of retirement, and would cost lives. Raising the quality of schools would raise the return to education, the inequality of earnings, and the risk of death. Better teachers are a health hazard. In the next section, I shall present some evidence, but first it is worth considering a simple but plausible model in which inequality is directly related to individual health.

One of the difficulties with national income inequality as a cause of individual mortality is the requirement of “action at a distance,” and the absence of any obvious mechanism through which it can operate. But local action is a different matter, especially if we switch the emphasis from inequality to relative income or rank; it is quite plausible that individual health should depend on the individual’s rank or status within the group to which that individual relates. Academics should have little difficulty with the notion that their health might depend on their salaries relative to those of their departmental colleagues. Studies of resilience in the face of disease and adverse shocks often identify the importance of high social status, if not in employment, then in another part of life, as when a janitor is an elder in his church. There is also a growing experimental and biomedical literature for both animals and humans that demonstrates the protective effects of rank on health, see in particular Robert Sapolsky (1993) and Sheldon Cohen (1991, 1997). This literature is beginning to put together a biochemical story for these effects which is currently the leading explanation for the link between socioeconomic status and health. But why should the existence of the gradient lead to an effect of inequality?

In human populations, we usually do not observe the reference group to which the individual belongs so we cannot make the ideal calculation relating individual health to individual income

relative to other members of the reference group. Instead, we observe only individual health and individual income. But individual income is partly informative about relative income. Someone with a very high income is more likely to have high income within his or her group, though it is always possible that he or she is a poor member of a very rich club; we have hit upon the poorest Wall Street speculator in the firm. Exactly how much individual income tells us about relative standing depends on the relative sizes of the variances within and between groups. If reference groups have little internal inequality, but there is great disparity between groups, individual income tells us little about relative income. Conversely, if the major source of inequality is within reference groups, which are themselves large and diverse, income tells us much more about relative income.

To formalize, consider the simple model

$$H = \alpha + \beta (y - \mu) \quad (3)$$

where H is some suitable measure of individual health, y is the (logarithm of) income, μ is mean income of the reference group to which the individual belongs, and α and $\beta > 0$ are parameters. The mean income of the group μ is distributed across the population of groups with mean $\bar{\mu}$, so that taking expectations unconditional on reference group, we have

$$E(y) = E(\mu) = \bar{\mu}. \quad (4)$$

We do not observe μ , only y , so that the best that we can do is to examine the expectation of H conditional on y ,

$$E(H|y) = \alpha + \beta (y - E(\mu|y)). \quad (5)$$

To calculate the expectation on the right-hand side, we need to make assumptions about the distributions of both. Although the argument can be made much more general, I work with the

simplest possible case. Suppose that, within the group, individual income is normally distributed around the group mean

$$y|\mu \sim N(\mu, \sigma_\varepsilon^2) \quad (6)$$

so that σ_ε^2 is the within-group variance, which serves as my measure of within group inequality. (Note that with y as the logarithm of income, neither (6) nor the following (7) are unrealistic.)

Group mean incomes are themselves normally distributed around their mean

$$\mu \sim N(\bar{\mu}, \sigma_\mu^2) \quad (7)$$

with σ_μ^2 the between-group variance.

Because (6) and (7) ensure that y and μ are jointly normally distributed, the expectation of μ conditional on y is a linear function of y and can be calculated from the least squares projection of μ on y . Write the conditional expectation as

$$E(\mu|y) = a + by \quad (8)$$

for some a and b , and note that we can always write

$$\mu = y - \varepsilon, \quad (9)$$

where ε is the distance of y from the group mean. Then we have

$$b = \frac{\text{cov}(y, \mu)}{\text{var } y} = \frac{\sigma_\mu^2}{\sigma_\varepsilon^2 + \sigma_\mu^2} \quad (10)$$

$$a = \bar{\mu} - b\bar{\mu} = \frac{\sigma_\varepsilon^2 \bar{\mu}}{\sigma_\varepsilon^2 + \sigma_\mu^2} \quad (11)$$

so that, substituting into (8), we have

$$E(\mu|y) = \bar{\mu} + \frac{\sigma_\mu^2}{\sigma_\varepsilon^2 + \sigma_\mu^2} (y - \bar{\mu}). \quad (12)$$

This equation is readily interpreted. If we know nothing about an individual, our best guess of his or her reference group mean income is the grand mean $\bar{\mu}$. But once we observe income, we can do better, adjusting our expectation up if income is greater than the grand mean, and down if it is below. The extent of the adjustment depends on how informative y is about μ , which depends on the two variances or, more precisely, on their ratio. Another useful interpretation comes from thinking about income y as an error-ridden estimate of the group mean μ . If we had data on μ , and we regressed μ on itself, we would get a slope of one and an intercept of 0. But if we regress μ on y , which is an imperfect measure of μ , the coefficient is asymptotically biased down by the familiar attenuation bias of least-squares in the presence of measurement error. (Note however that this argument rests on the linearity of the regression function (8) and that, without joint normality, the conditional expectation of μ will typically not be linear in y .)

As a final step, combine equation (12) with equation (5) to give the relationship between the expectation of individual health and individual income; the regression function is

$$E(H|y) = \alpha + \frac{\beta \sigma_{\epsilon}^2}{\sigma_{\epsilon}^2 + \sigma_{\mu}^2} (y - \bar{\mu}). \quad (13)$$

Equation (13) formalizes the initial argument. When within group variance is large relative to between group variance, so that σ_{ϵ}^2 is large relative to σ_{μ}^2 , income is a good indicator of relative income and is a strong predictor of health. In the opposite situation, when the groups are much more homogeneous than the population as a whole, σ_{μ}^2 is much larger than σ_{ϵ}^2 , and there is little or no relationship between health and individual income. The model therefore has predictions about how the gradient should change as we work with different comparison groups. For example, if we had data on a firm or on a single employer—the Whitehall study is perhaps the most

obvious example—the gradient should be steeper than the gradient between health and income in the population as a whole, or between populations in international comparisons. If people tend to sort themselves into reference groups of people with rather similar incomes, then we would expect to see very little gradient in the population as a whole.

As interesting as the implications of the model for the gradient is its prediction that, at a given income level, the ratio of within- to between-group inequality has an effect on individual health. This is true in spite of the fact that inequality has no explicit role in the determination of health in equation (3); note for example that spreading within-group incomes away from the group mean will raise within-group inequality but have no effect on average group health. The effect works by rotating the relationship between health and income around the point of mean income and mean health. An increase in overall income inequality will have different effects depending on how it breaks down into between and within group components. If within group inequality is not affected, as might be the case if wage structures within industries of firms stay fixed but spread out across firms then, conditional on income, individual health will decline for those above the mean and increase for those below the mean. If the inequality increase comes from a within-group increase, health will improve for the rich and decline for the poor. And if inequality increases happen equally within groups and between them, the relationship between health and income will be undisturbed. Since we know very little about the relevant reference groups for these purposes, all of these possibilities are worth serious investigation, as is the general proposition that inequality changes work through a rotation of the gradient.

It is important to note that, in spite of the relationship between health, income, and inequality in (13), changes in inequality have no effect on anyone's health except to the extent that they

change people's relative incomes within their reference groups. To see the point, suppose that the structure of reference groups is such that we can design a change in tax policy that reduces the *between-group* inequality σ_{μ}^2 without affecting *within-group* inequality σ_{ϵ}^2 . By (13), the policy change will increase the average health of people within income groups above average income, and decrease the average health of within income groups below the average. But the change in inequality has no effect on anyone's health. Health is determined by relative income in the reference group which, by assumption, is unchanged. But because the relative variances have changed, there has been a change in the composition of (reference group) relative incomes among those in any given (absolute) income group, which is why average health changes conditional on income. Just because inequality plays a role in the regression of health on income does not imply that policy-induced changes in inequality will have any effect on average health.

This simple model of health and relative income can be straightforwardly modified to permit a more fundamental role for within-group inequality. Suppose that instead of making health depend on relative income within the reference group, we make health a declining function of the distance of each individual from the top-ranked member of the group. In a linear form, this can be written

$$H = \tilde{\alpha} - \beta(y_{\max}^{\mu} - y) \quad (14)$$

so that $\tilde{\alpha}$ is the health of the best-off member of the group and β measures the effect of having less than the best-off member of the group and is, as it was in (3), the slope of the within-group gradient. Indeed, in a single cross section, (14) is indistinguishable from (3). For simplicity, and to relate this model to the earlier one, suppose that, within each group, the top person's income is θ standard deviations above the mean, so that

$$y_{\max}^{\mu} = \mu + \theta \sigma_{\varepsilon}. \quad (15)$$

Combining (14) with (15) gives a health equation that is a straightforward modification of the equation (3) with which we began, i.e.

$$H = \tilde{\alpha} + \beta(y - \mu) - \beta\theta\sigma_{\varepsilon}. \quad (16)$$

This formulation seems more in the spirit of Wilkinson's book in that inequality has a direct negative effect on individual health. In the original, purely relative formulation (3), an increase in within group spread had no effect on average group health, because the decrease in health of those who became worse-off was offset by the increase in the health of the more fortunate members of the group. But in (15), it is the income gap from the top that matters, so that the health of the top person is not improved if his income rises, while that of everyone below him is hurt by the widening gap from the top. As a result, inequality has a direct effect on health of each member of the group as well as on the average health of the group. Note also that the choice of the top person as reference is entirely arbitrary, and thus unsatisfactory without further argument. If instead of (15), the health of the *bottom* person was made unresponsive to income, with the health of others increasing with the income gap from the bottom, increased income inequality would *increase* health on average, and the last term in (16) would appear with a positive rather than a negative sign. The model is therefore silent on why inequality should be a health hazard, though it provides a usual framework for empirical analysis. Together with the original model, it also provides a useful counterexample to a proposition that is sometimes asserted in the literature, that the dependence of health on relative income implies that average health is negatively linked to inequality. It is possible to write down plausible models of individual health and relative income which generate positive, negative, or no relationship between inequality and average health.

In the absence of information on the groups, the new formulation is no more directly useful than was the original. But the same arguments used for the original model can be used for the new model to derive an observable relationship between health and income. Since the calculations are identical, I simply record the final result. Instead of equation (13), there is an additional term in inequality and we now have

$$E(H|y) = \alpha + \frac{\beta\sigma_{\varepsilon}^2}{\sigma_{\varepsilon}^2 + \sigma_{\mu}^2}(y - \bar{\mu}) - \beta\theta\sigma_{\varepsilon}. \quad (17)$$

Within group inequality now plays two roles, its original one of rotating the gradient, and a new, more fundamental one, as a direct determinant of average health. Once again, it is not income inequality for the economy as a whole that has the direct role, only inequality within each reference group.

3. Empirical evidence

Much of the evidence in favor of a link between inequality and growth is provided in Wilkinson's book and articles, in particular Wilkinson (1995, 1996). For the United States, George Kaplan et al (1996) and Bruce Kennedy et al (1996) have linked state-level mortality to state level income inequality, though their results, like Wilkinson's, have been challenged on a number of fronts, including the choice of inequality measure, lack of robustness to inclusion or exclusion of alternative plausible controls—such as fraction black—and aggregation fallacies, especially the failure to allow for the effect of inequality that results from the aggregation of a nonlinear relationship between income and mortality at the individual level, Jennifer Mellor and Jeffrey Milyo (1999).

Some of the most spectacular evidence is given in Figures 3, 5, and 6 in Wilkinson (1994). The first of these scatter diagrams shows no relationship across OECD countries between the changes from 1970 to 1990 in life expectancy and GDP per capita in PPP \$. The second graph, for about half of the countries in the first, shows life expectancy in 1970 and the gini coefficient of post-tax income inequality, standardized for household size, and has a correlation coefficient of -0.81 . The final graph, for 12 EEC countries, shows the relationship between the annual rates of change from 1975 to 1985 in life-expectancy and the percentage of the population in relative poverty, with a correlation of -0.73 . The problem with such comparisons, of course, is the quality and difficulty of interpretation of the distributional data, even among OECD countries, see Ken Judge (1995) and Wilkinson (1995) for an exchange over this issue (among others.) While I suspect that the hypothesis is probably not ultimately resolvable using cross-country comparisons, or at least not without detailed examinations of and adjustments to the survey data for each country, it is useful to begin by reporting these data.

I have taken international data on gini coefficients of the distribution of household income from Deininger and Squire (1994); I have used data from “industrialized” countries and included only those measurements whose quality is graded “acceptable” by Deininger and Squire, excluding measurements based on household expenditures rather than income. Data on life expectancy are taken from the 1998 edition of the World Bank’s *World Development Indicators*. In the top left panel of Figure 1, I show those observations available for any year in the interval 1980 to 1982, in most cases with two observations for each country. In the top right panel, the exercise is repeated for data for the years 1987 to 1990. The dates are chosen so as to maximize the number of countries for which data are available. In the left-hand graph, the simple correlation coefficient

is -0.26 , which has the right sign, though the scatter is wide. In the right hand panel, on average seven to eight years later, life expectancy is higher—the points are shifted upwards—but the correlation has fallen to only -0.14 . Note that the United States and Japan are reported as having high income inequality in both periods, while the United Kingdom, which starts out as one of the most equal countries in the earlier picture, is in the middle of the rankings in the later one. The bottom-left panel shows the relationship between the changes in life expectancy and the changes in the gini coefficients. The points shown are the differences in the averages (in the case of multiple observations) in the top two figures, and excluding Denmark, which appears only in the later period. Here the negative correlation has become a positive correlation. In these data, there is no clear and consistent correlation between inequality and life expectancy.

Three countries, the Japan, the United Kingdom, and the United States, have enough data to allow tracking of both life expectancy and the gini coefficient on a consistent basis over a run of years. The bottom right panel shows what happened for each country for some years from 1970 to 1990 (1991 for the United Kingdom.) In all three countries, life expectancy rose over these two decades, and it rose by the largest amount in Japan, which had no increase in income inequality. In the United States, life expectancy also rose, more rapidly in the 1970s than in the 1980s (a fact to which we shall return), and in total over the two decades by almost the same amount as the rise in the United Kingdom. But the increase in income inequality in the United Kingdom was about twice as large as the increase in the United States. With enough imagination, proponents and opponents of the thesis could each find something in these figures but the immediate visual impression is of a simultaneous increase in income inequality and life-expectancy.

Figure 2 shows more detail on the trends in males and female mortality in the United States;

these data are taken, as are those used below, from the Berkeley Mortality Database (BMD) website at <http://demog.berkeley.edu/wilmoth/mortality>. I have used the estimated US population by sex in the fourth quarter of 1984 (from the Bureau of the Census website) to weight the mortality rates. The rapid decline in mortality from about the mid-1960s (earlier for women than for men) was much slower after about 1980, a phenomenon that could come from many causes, including the increase in inequality. The decrease in female relative to male mortality, which is apparent in the figure until the mid-1960s, is the end of a long secular trend in favor of females, mostly associated with the relative reduction in mortality among young women compared with young men.

To investigate more closely the relationship between mortality and inequality, I have combined the mortality data from the BMD with data on household income and household income inequality from the March Current Population Surveys (CPS) from 1981 to 1993. My basic unit of analysis is a birth cohort; defined as the group of people born in a single year, and I examine the relationship between mortality for that cohort in any particular calendar year in relation to the level and inequality of income of that cohort. Note that the March CPS collects data on income for the *previous* year, so that by matching mortality and CPS data for the same date, I am matching deaths in 1985(say) with incomes and income inequality in 1984. A one year lag is a reasonable specification with which to start.

Note that I am not claiming that cohorts are the relevant reference groups within which health is determined as in the theory of Section 2. But birth cohorts are likely to contain a higher ratio of relevant to irrelevant reference people, which should steepen the gradient and make it easier to detect than is the case in the aggregate or international data. The use of cohort data also solves

the scarcity of microeconomic data containing full economic information on decedents, and allows the combination of different sources, here the CPS for economic information, and vital statistics for mortality. Cohorts also have decisive advantages over aggregate data in terms of sample size. With an age restriction to those between 25 and 75 inclusive, the 13 years of CPS data yield 663 cohort-year observations, from those born between 1906 (observed at age 75 in 1981) and 1968 (observed at age 25 in 1993.)

The cohorts in the CPS are defined by the age (or date of birth) of the *household head*, (or more accurately, the *reference person*) and income and income inequality are represented by the cohort mean of the logarithms of income and the cohort variance of the logarithm of income, again calculated on a *household* basis. When I merge these data with the mortality data, I am merging in the mortality rates for *individuals*. While the conjunction of households and individuals is uncomfortable, particularly for females who are less likely to be household heads, there is no ideal method of reconciling the household data in the CPS with death rates of individuals, and the difficulty would exist even with unit record data. Income is observed for households, not individuals. Even so, it would be useful to redo the calculations imputing household incomes to individuals, and merging on individual birth cohorts, something that would also allow me to look at individual earnings and other components of income. I hope to report these results in subsequent versions of this paper. My use of the variance of logarithms as the measure of inequality should also be seen as a first cut. Trends in income inequality in the U.S. are not invariant to the selection of inequality measure; the variance of logarithms is (too) sensitive to developments at the bottom of the distribution, and thus to imputation procedures for zero or non-reporting, while another obvious measure, the gini, is (too) sensitive to the top of the distribution, and thus to

details of top-coding. Again, more work remains to be done.

There are a number of important issues that have to be borne in mind when thinking about the results. One is that I have made no attempt to correct the estimates for the attenuation bias to the coefficients that comes from the fact that the cohort income and inequality estimates are subject to sampling errors. A correction is feasible following the methods in Deaton (1985); for the moment the important thing to note is that the estimates given here are conservative, in the sense of being biased toward zero and insignificance.

Perhaps the most central issue is the dependence of mortality rates on age, and how this dependence is taken into account in testing the hypotheses. Between ages 25 and 75, and except possibly for the first two or three years, the probability of dying rises monotonically with age. So does the variance of household income, see Deaton and Paxson (1994) for evidence and reasons. But the hypothesis is presumably not that inequality rather than biology is the fundamental determinant of aging and death. There are various ways of adjusting mortality rates for age, and the effects of some of these are explored in the tables. It should be noted that age profiles of mortality have not been constant over time in the US, for either males or females. The biggest decreases in female mortality rates have been at younger ages, and in male mortality rates at older ages, so that the female age-profile has become steeper and the male age-profile flatter.

The entries at the top left of Table 1 show what happens with the simplest form of age adjustment, which is the inclusion of age and its square in the log odds regression. These baseline regressions, run on 663 cohort-year observations and separately for men and women, have the log odds of dying on the left-hand side and age, age squared, and the mean and variance of the logarithms of household income on the right hand side. They show a significant negative effect of

income and a significant positive effect of the variance, as hypothesized. Income is more protective for men than for women; the male coefficient is twice as large. It is tempting to attribute this to the matching procedure, which by matching individual age to the age of the household head, typically serves men better than women, but even in studies with unit record data, the mortality gradient with household income tends to be steeper for males than for females, see e.g. Rogot et al (1992). The effect of income is large; doubling household income would reduce a man's risk of death by about a third, about two thirds of the way to his wife's level. The estimated effect of inequality is absolutely smaller. Between 1981 and 1993, the average over cohorts of the variance of log income rose from 0.67 to 0.82, so that according to the estimates, which are very similar for men and women, the average log odds of dying should have risen by 0.026. But this is quantitatively important relative to the reduction in the cohort average of the log odds, which from 1981 to 1995 was 0.07 for both men and women, so that this estimate attributes about half of the slowdown in the rate of mortality decline to the rise in income inequality.

However, the estimated effects of inequality are not robust to the way in which the mortality schedule is age standardized. In the second pair of columns, I report what happens if, instead of using the log odds of dying as the left-hand side variable, I subtract out the estimated age effects from a regression of the log odds on an unrestricted set of age and cohort dummies. In order to obtain maximum accuracy, this preliminary regression is estimated on (almost) the full set of data in the BMD, i.e. on all ages between 0 and 85 for all years from 1900 to 1995. The unrestricted age effects allow for any shape in the age mortality profile, and the unrestricted cohort effects allow for parallel shifts down (or up) in the age-profiles across cohorts. I included the

cohort effects in order to limit possible bias to the age effects if they were omitted, but the results in the tables are almost identical if the cohort effects are excluded, and the adjustment made by using the residuals from a regression of the log odds on unrestricted age effects. Even after adjustment, I still include age and age squared in the regressions to capture age effects that remain. Cohort age-profiles of the log odds of mortality are not perfectly described by the sum of age and cohort effects—though the approximation is quite good after age 25, especially for men.

This superior age-adjustment procedure leaves the estimated effects of income more or less unaltered, but eliminates the effect of inequality for men, and reduces it for women. Very similar results are shown in the next pair of estimates in the same row, where I adjust the log odds of dying by subtracting the age-specific log odds of dying in 1940, a simpler procedure which gives results very close to those obtained by first eliminating age and cohort (or just age) effects.

The last pair of estimates in the first row shows what happens when I use the raw log odds (or indeed any of the adjusted measures) and include in the regression an unrestricted set of age dummies. Once again, the highly significant and (somewhat smaller but still large) protective effect of income remains, but the signs on the variance terms have changed, and increases in inequality are estimated to be protective, albeit with a small coefficient. In this first set of regressions, it appears to be the positive relationship between income inequality and age that gives an apparent positive effect of inequality on mortality. When age is controlled for in a flexible way, the effect either vanishes or is reversed.

One argument against this first set of models is that inequality and income take time to work themselves through into mortality, and that we should not expect to see a relationship on a year to year basis. To test this rigorously would require more information on how long the process is

supposed to take; for example, it would be possible to test for lagged effects, or a hypothesis that inequality is most harmful at specific ages. One plausible hypothesis is that it is lifetime exposure to income and income inequality that determines mortality. This hypothesis can be tested by projecting the average logarithm of income and the variance of income on a set of cohort dummies and including only the projections in the regression. In the case of the mean logarithm of income, this is the same decomposition that is used in life-cycle models of saving to recover estimates of lifetime resources, or permanent income. In practice, it is convenient to run these regressions as before, but using unrestricted cohort dummies as instrumental variables for income and income inequality. In the second panel of Table 1, I instrument only for the mean of log income, so that income works through permanent income, and inequality works instantaneously. In the bottom panel, I instrument for both income and inequality, so that both work only over the lifetime.

The most important result in the middle panel—repeated in the bottom panel—is that permanent income is more protective than current income, twice or three times as much. If income is protective at all, it is perhaps not surprising that permanent income is more protective, or indeed that current income appears only as a proxy for permanent income. Note that, as before, the effects are smaller for females than males, but the larger effect of permanent over current income appears for females just as for males.

Instrumenting for permanent income has relatively little effect on the inequality estimates, when inequality itself is not instrumented. Inequality is significant when mortality is not fully adjusted for age, but its effect becomes small, insignificant, or perverse in the later columns. However, the situation is rather different in the final row, where both income and inequality are

instrumented. Permanent inequality is estimated to raise mortality, and to do so by much more than when we use only contemporaneous inequality. These effects, although reduced by better age adjustment in the second and third sets of estimates, remain quantitatively large and statistically significant until the last pair of columns, where age effects are fully controlled, at which point the effects of inequality vanish. Once again, it is the age structure of inequality that is generating the estimated effects on mortality, which disappear once inequality is prevented from working through age by the inclusion of a full set of age dummies.

If we had prior knowledge of the non-economic determinants of the age-profile of mortality, and could thus perfectly adjust the age profile, it would be possible to investigate whether income inequality over the life-cycle plays a role in shaping the age-profile of mortality. But absent that knowledge, and given the strong (nonlinear) relationships between age on the one hand and both mortality and inequality on the other, it is necessary to include the full set of age dummies, and to conclude that there is no evidence in these data that higher inequality raises the risk of dying. The lack of robustness of the effects of inequality on mortality, and in particular the sensitivity to methods of age-adjustment, should also sound a note of caution for other attempts to study this mechanism.

Table 2 contains an important set of cross-checks. Because both the level of income and its dispersion are rising over time, and because mortality is declining over time, there is some risk that the results might be affected by a spurious correlation whereby one or other of income and its inequality picks up the trend in mortality. A simple way of testing for this possibility is to include a time trend in all the regressions. But the results in Table 2, although differing in detail, do not differ in their broad import. In the preferred specification, with both variables instrumented

and a full set of age dummies, the effect of mean income on male mortality is larger than before, and that on female mortality smaller. Inequality has a negative (i.e. perverse) but insignificant effect on male mortality, and a positive effect on female mortality.

Finally, I turn to the theoretical model developed in Section 2 and its prediction that inequality should change the slope of the gradient. Testing this hypothesis requires entering the product of the mean of log income and the variance of log income in the regressions. In most of the regressions in Tables 1 and 2, it is not possible to reject the hypothesis that the interaction is zero, a partial exception being an estimated positive interaction term for females in the regression in the first column of the third row, where both the variance and the mean (as well as their interaction) are instrumented, but where the mortality schedule is unadjusted except for the inclusion of age and age squared in the regression. Since the age-adjustment here is almost certainly inadequate, this result is of little consequence.

More interesting is what happens in the preferred specification in the last cell of the table, where both income and income inequality are instrumented by cohort dummies and where the regression includes a full set of age dummies. Table 3 shows the three coefficients for males and for females, with and without time trends. It also shows the derivatives of the log odds of mortality with respect to the mean and the variance calculated at the mean values of both. As always, income exerts a strong protective effect and, as in the original preferred specification, inequality has no effect at the mean. But the interaction terms are highly significant and always negative so that, according to these estimates, increases in inequality rotate the mortality to income gradient around the mean, with higher inequality steepening the gradient. This is what the reference group-theory predicts in the model without a fundamental effect of inequality and provided

inequality acts in such a way as to increase within-group inequality relative to between group inequality. Of course, there is no independent evidence that inequality works in this way, but the results are certainly consistent with that version of the theory.

4. Conclusions and caveats

This paper has proposed a simple model in which an individual's health is affected by the individual's income relative to the average income of members of a reference group. Because reference groups can be defined in myriad ways, reference group incomes are hard to observe. Nevertheless, individual income is informative about reference group income, to an extent that depends on the ratio of between-group and within-group variances of income, with the result that the "gradient," the slope of the expectation of health conditional on income, depends on the ratio of the variances. In this way, inequality affects the relationship between income and health. My preliminary examination of mortality in the United States is consistent with this story. Increasing inequality does not appear to have a direct influence on mortality, but the protective effect of income against mortality has become stronger as income inequality has increased. According to the theory, these results do not imply that policy changes that affect inequality will necessarily have any effect on aggregate (or even individual) mortality. When the ratio of between-group to within-group inequality changes, any given income group has a different mix of high status and low status individuals, and it is the change in this mix that changes health conditional on income, not any change in health.

While I hope that these theoretical and empirical findings will be a useful addition to the debate on socio-economic differences in health, I wish to conclude by referring to some of the

important issues that I have ignored. First, the model developed here makes no concession to the other direction of causality, from health to income. This “reverse” link is undoubtedly present, but it is not the topic of this paper. Second, in order to focus on relative income, I have ignored other determinants of health. Some researchers believe that the effect of income on health is the result of the association between income and education, and while my own view is that both education and income are separately protective, I do not wish to downplay the independent importance of education. Again, I have ignored it only to focus on other matters. Finally, by working with linear models, I have ignored one of the leading candidates for explaining the effects of inequality on average health, which is aggregation over a concave relationship between health and income. While such a mechanism is plausible, the lack of a link between aggregate mortality and inequality in my data could be interpreted as a defense of linearity.

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.
- Allison, R. Andrew and James E. Foster, 1998, "Measuring health inequality using qualitative data," Vanderbilt University, processed.
- Atkinson, Anthony B., 1970, "On the measurement of inequality," *Journal of Economic Theory*, **2**, 244–63.
- 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 and Christina H. Paxson, 1994, "Intertemporal choice and inequality," *Journal of Political Economy*, **102**, 437–67.
- Deaton, Angus and Christina H. Paxson, 1998a, "Aging and inequality in income and health," *American Economic Review, papers and proceedings*, **88**, 248–53.
- Deaton, Angus and Christina H. Paxson, 1998b, "Health, income and inequality over the life-cycle," in David A. Wise, ed., *Frontiers in the Economics of Aging*, Chicago. Chicago University Press for the NBER, 431–57.
- Deininger, Klaus and Lyn Squire, 1996, "A new data set measuring income inequality," *World Bank Economic Review*, **10**, 565–91.
- Elo, Irma T., and Samuel H. Preston, 1996, "Educational differentials in mortality: United States, 1979–85," *Social Science and Medicine*, **42**, 47–57.
- Feldman, Jacob J., Diane M. Makuc, Joel C. Kleinman, and Joan Cornoni-Huntley, 1989, "National trends in educational differentials in mortality," *American Journal of Epidemiology*, **129**, 919–33.
- Judge, Ken, 1995, "Income distribution and life expectancy: a critical appraisal," *British Medical Journal*, **311**, 1282–5
- 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.
- Kitigawa, Evelyn M. and Philip M. Hauser, 1973, *Differential mortality in the United States: a study in socioeconomic epidemiology*, Cambridge, MA. Harvard.
- Le Grand, Julian, 1987, "Inequalities in health: some international comparisons," *European Economic Review*, **31**, 182–91.
- Lynch, J. W., G. A. Kaplan and J. T. Salonen, 1997, "Why do poor people behave poorly?" *Social Science and Medicine*, **44**, 809–31.
- Macintyre, Sally, 1997, "The Black report and beyond: what are the issues?" *Social Science and*

- Medicine*, **44**, 723–45.
- Meara, Ellen, 1998, “Why is health related to socioeconomic status?” Department of Economics, Harvard University, processed.
- Mellor, Jennifer M. and Jeffrey Milyo, 1999, “Income inequality and health status in the United States: Evidence from the current population survey,” Departments of Economics, The College of William and Mary and Tufts University, processed.
- Pappas, G., S. Queen, W. Hadden and G. Fisher, 1993, “The increasing disparity in mortality between socioeconomic groups in the United States, 1960 and 1986,” *New England Journal of Medicine*, **329**, 103–15.
- Preston, Samuel H. and Irma T. Elo, 1995, “Are educational differentials in adult mortality increasing in the United States?” *Journal of Aging and Health*, **7**, 476–96.
- 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.
- Sapolsky, Robert M., 1993, “Endocrinology alfresco: psychoendocrine studies of wild baboons,” *Recent Progress in Hormone Research*, **48**, 437–68.
- Sen, Amartya K., 1973, *On economic inequality*, Oxford. Clarendon Press.
- Wilkinson, Richard G., 1994, “The epidemiological transition: from material scarcity to social disadvantage,” *Daedalus*, **123**, 61–77.
- Wilkinson, Richard G., 1995, “Commentary: a reply to Ken Judge: mistaken criticisms ignore overwhelming evidence,” *British Medical Journal*, **311**, 1285–7.
- Wilkinson, Richard G., 1996, *Unhealthy societies: the afflictions of inequality*, London, Routledge.

Table 1

Effects of mean log income and the variance of log income on the log odds of dying: birth cohorts observed in 1981–93

	Age unadjusted		Age adjusted		Adjusted relative to 1940		Including full age effects	
	males	females	males	females	males	females	males	females
OLS								
mean lny	-0.45 (8.9)	-0.22 (6.0)	-0.42 (10.)	-0.27 (8.1)	-0.54 (13.)	-0.29 (8.0)	-0.35 (7.6)	-0.28 (11.)
var lny	0.17 (6.3)	0.16 (7.9)	0.02 (1.0)	0.12 (6.7)	0.02 (1.1)	0.15 (7.8)	-0.09 (3.7)	-0.06 (4.5)
IV for m(lny)								
mean lny	-1.15 (13.)	-0.62 (10.)	-1.04 (15.)	-0.64 (12.)	-1.15 (16.)	-0.67 (12.)	-0.73 (11.)	-0.53 (14.)
var lny	0.12 (4.0)	0.13 (6.0)	-0.02 (0.7)	0.09 (4.8)	-0.02 (0.6)	0.12 (5.8)	-0.12 (4.6)	-0.08 (5.4)
IV for both								
mean lny	-1.38 (12.)	-0.81 (9.3)	-1.18 (14.)	-0.81 (10.)	-1.30 (14.)	-0.80 (10.)	-0.82 (10.)	-0.62 (12.)
var lny	0.71 (9.3)	0.59 (11.)	0.32 (5.8)	0.52 (10.)	0.35 (6.2)	0.58 (11.)	0.03 (0.4)	0.06 (1.3)

Notes: *t*-values in parentheses. The dependent variable is the log odds of dying for each birth cohort in each of the 13 years, 1981 to 1993. Age and age squared are included in all regressions except those with a full set of age effects. In the “age unadjusted” regressions, the dependent variable is the unadjusted log odds. In the “age adjusted” regressions, the dependent variable is log odds less the estimated age-effects from a regression of the log odds on unrestricted age and birth cohort dummies, calculated for ages 0 to 85 for all years from 1900 to 1995. “Adjusted relative to 1940” means that the dependent variable is the log odds of dying minus the log odds of dying in 1940. In the final columns, the dependent variable is the unadjusted log odds and an unrestricted set of age dummies are included in the regression. The instrumental variables are cohort (date of birth) dummies, so that the IV procedure is equivalent to projecting the mean or variance (or both) on lifetime cohort effects prior to inclusion in the regression.

Table 2

Effects of mean log income and the variance of log income on the log odds of dying: birth cohorts observed in 1981–93
(with a time trend included in all regressions)

	Age unadjusted		Age and cohort adjusted		Adjusted relative to 1940		Including full age effects	
	males	females	males	females	males	females	males	females
OLS								
mean lny	-0.39 (6.8)	-0.05 (1.2)	-0.39 (8.4)	-0.12 (3.3)	-0.55 (11.)	-0.14 (3.5)	-0.30 (5.2)	-0.07 (2.8)
var lny	0.18 (6.6)	0.20 (11.)	0.03 (1.5)	0.15 (8.9)	0.02 (1.1)	0.19 (10.)	-0.07 (2.8)	-0.01 (0.6)
IV for m(lny)								
mean lny	-1.70 (12.)	-0.43 (5.2)	-1.61 (13.)	-0.50 (6.7)	-1.82 (14.)	-0.56 (6.9)	-1.13 (10.)	-0.39 (6.8)
var lny	0.03 (0.8)	0.16 (7.1)	-0.11 (3.2)	0.11 (5.7)	-0.13 (3.5)	0.14 (6.4)	-0.22 (6.2)	-0.05 (3.3)
IV for both								
mean lny	-1.68 (10.)	-0.41 (3.4)	-1.60 (12.)	-0.48 (4.4)	-1.81 (13.)	-0.53 (4.5)	-1.24 (10.)	-0.44 (6.4)
var lny	0.62 (7.1)	0.72 (12.)	0.20 (3.0)	0.61 (11.)	0.20 (2.9)	0.67 (12.)	-0.18 (1.9)	0.15 (2.9)

Notes: *t*-values in parentheses. The dependent variable is the log odds of dying for each birth cohort in each of the 13 years, 1981 to 1993. Age and age squared are included in all regressions except those with a full set of age effects. In the “age unadjusted” regressions, the dependent variable is the unadjusted log odds. In the “age adjusted” regressions, the dependent variable is log odds less the estimated age-effects from a regression of the log odds on unrestricted age and birth cohort dummies, calculated for ages 0 to 85 for all years from 1900 to 1995. “Adjusted relative to 1940” means that the dependent variable is the log odds of dying minus the log odds of dying in 1940. In the final columns, the dependent variable is the unadjusted log odds and an unrestricted set of age dummies are included in the regression. The instrumental variables are cohort (date of birth) dummies, so that the IV procedure is equivalent to projecting the mean or variance (or both) on lifetime cohort effects prior to inclusion in the regression.

Table 3
Interactions between means and variances in mortality regression with full age effects

	Males		Females	
<i>With no time trend</i>				
mean lny	-0.386	(2.3)	0.012	(0.1)
variance lny	7.309	(3.0)	10.58	(6.1)
mean * variance	-0.732	(3.0)	-1.058	(6.1)
$\partial/\partial m$ at mean	-0.951	(9.9)	-0.805	(12.)
$\partial/\partial v$ at mean	0.042	(0.5)	0.078	(1.3)
<i>With time trend</i>				
mean lny	-0.615	(2.9)	0.036	(0.3)
variance lny	12.68	(4.0)	10.02	(5.6)
mean * variance	-1.302	(4.1)	-0.998	(5.5)
$\partial/\partial m$ at mean	-1.621	(9.3)	-0.735	(7.5)
$\partial/\partial v$ at mean	-0.236	(2.1)	0.107	(1.7)

Notes: *t*-values in parentheses. The means of mean(lny) and var(lny) are 9.93 and 0.77 respectively, and the last two rows in each panel show the derivatives evaluated at those values. Except for the inclusion of the interaction terms, these regressions correspond to those shown in the last columns of Tables 1 and 2.

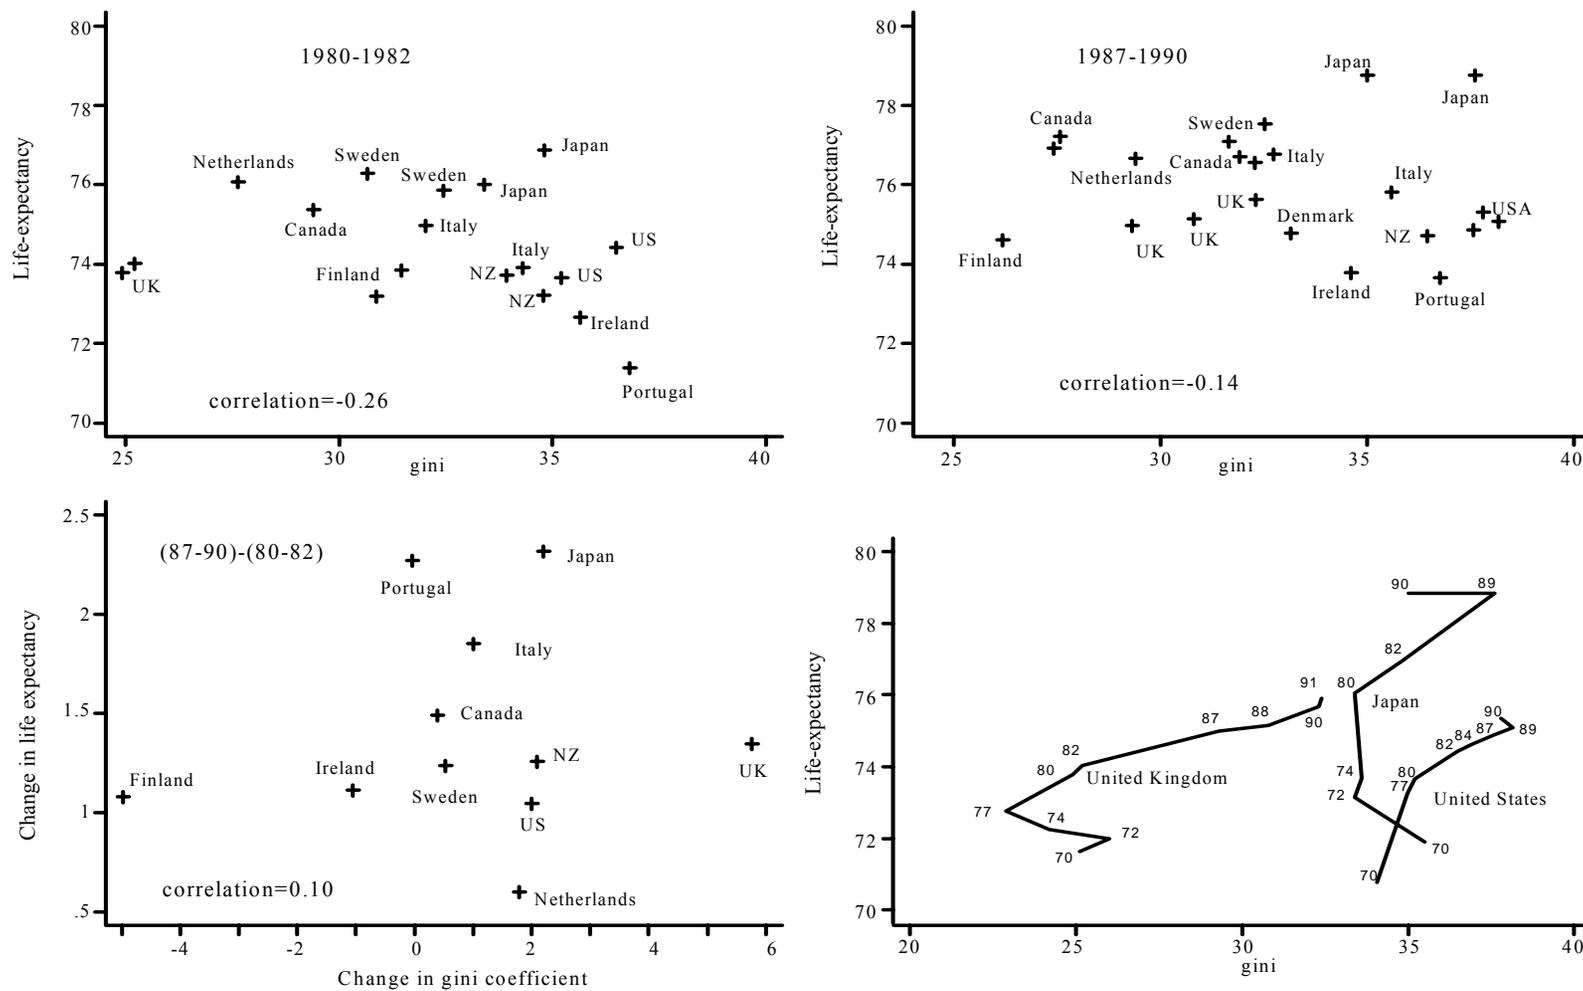


Figure 1: Inequality and life expectancy: some international comparisons

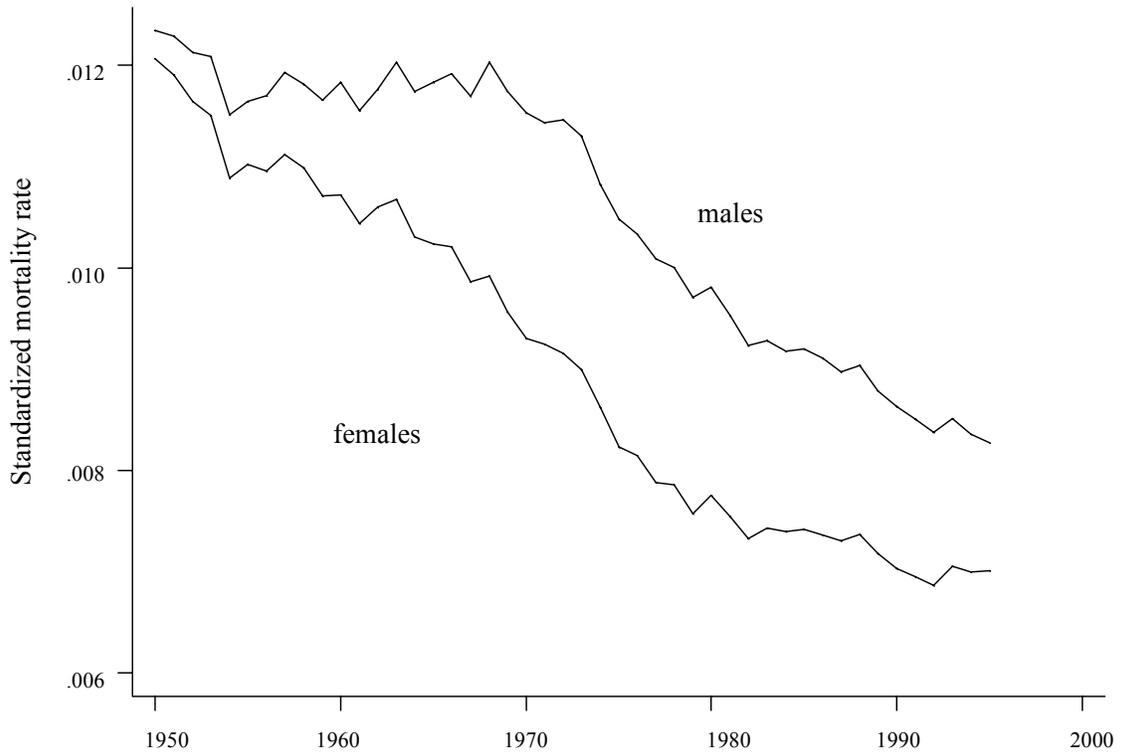


Figure 2: Standardized mortality rates for males and females: United States, 1950–1995