

This PDF is a selection from an out-of-print volume from the National Bureau of Economic Research

Volume Title: Frontiers in the Economics of Aging

Volume Author/Editor: David A. Wise, editor

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-90304-4

Volume URL: <http://www.nber.org/books/wise98-1>

Publication Date: January 1998

Chapter Title: Health, Income, and Inequality over the Life Cycle

Chapter Author: Angus S. Deaton, Christina Paxson

Chapter URL: <http://www.nber.org/chapters/c7309>

Chapter pages in book: (p. 431 - 462)

10 Health, Income, and Inequality over the Life Cycle

Angus Deaton and Christina Paxson

10.1 Introduction

In previous work, Deaton and Paxson (1994, 1997), we used data from the United States, Great Britain, Taiwan, and Thailand to document that inequality increases within cohorts with age, for consumption, income, and earnings. In this paper, we extend the analysis to two health-relevant measures, the body-mass index and self-reported health status. We use data on more than 500,000 adults in the United States to track birth cohorts over time and to document the evolution of the two measures with age, looking at both cohort means and within-cohort dispersion. We also consider the life cycle profile of dispersion in income and health jointly, presenting evidence separately for men and women, and for blacks and whites.

Our original work on consumption and income inequality was motivated by the prediction of the standard theory of autarkic intertemporal choice that within-cohort inequality in consumption and income (although not necessarily earnings) should increase with cohort age, at least up to the date of retirement. Although the theory has no immediate extension to processes other than income and consumption, there are a number of reasons to extend the analysis to health status.

First, we wish to investigate the generality of the proposition that dispersion

Angus Deaton is the William Church Osborn Professor of Public Affairs and professor of economics and international affairs at Princeton University and a research associate of the National Bureau of Economic Research. Christina Paxson is professor of economics and public affairs at Princeton University and a research associate of the National Bureau of Economic Research.

The authors are grateful to Sudhir Anand, David Meltzer, Amartya Sen, Burt Singer, and Jim Smith for helpful discussions during the preparation of this paper. Research assistance was ably provided by Shelley Clark. The authors gratefully acknowledge financial support from the National Institute on Aging and from the John D. and Catherine T. MacArthur Foundation. Errors, omissions, and incompleteness are the responsibility of Benjamin Paxson. Gabinet, born 25 March 1997.

increases with age. For the four countries where we have looked, it is true of income, consumption, and earnings. We are curious as to whether the proposition is true for other state variables, such as weight, the body-mass index, self-reported health status, dexterity, intelligence, or ability to complete specified tasks.

Second, while health status is interesting as an example, it is also important in its own right. Inequalities in income and consumption are of concern because they are important components of welfare. But as we move from a narrow, economic measure of well-being toward broader definitions, health status has the most immediate claim on our attention. Nor is health status independent of economic status; indeed, there is a well-documented but poorly understood "gradient" linking socioeconomic status to a wide range of health outcomes (Adler et al. 1994), as well as changes in wealth and changes in health among the elderly (Smith 1995). Income or its correlates (smoking, obesity, social status, and various types of behaviors) may directly affect health, and where health care is expensive, the ability to pay may give access to superior health services. There are also mechanisms that operate in the opposite direction; health status affects the ability to work and, among the elderly, the timing of retirement. There is also a literature linking health status to relative deprivation, or to the income distribution: Wilkinson goes so far as to claim that "mortality rates in the developed world are no longer related to per capita economic growth, but are related instead to the scale of income inequality in each society" (1994, 61). Yet there has been no research of which we are aware that tracks these relationships over the life cycle, or that looks at the life cycle patterns for clues to directions of causality.

Third, it is plausible that the theoretical reasons that consumption, income, and earnings processes disperse also apply to health status. This requires a little explanation. In Deaton and Paxson (1994), we started from the implication of (some) theories of intertemporal choice that individual consumption should follow a martingale process; Hall (1978) showed that under appropriate assumptions, consumption is the cumulative sum of uncorrelated increments. The same will be true of earnings if employers pay workers their expected marginal product (see Farber and Gibbons 1996). It is also plausible that, at least in part, health status should be a cumulative process, determined by the "piling up of adverse life experiences" (Singer and Ryff 1997) offset by recuperative processes. Although some health shocks will have only temporary effects, others will leave a permanent residue, so that even if this residue is a small component of the original shock, the resulting health status will be non-stationary. Because most of the effects of most health shocks wear off over time, there is no reason to suppose that health status is a martingale; on the contrary, health status of individuals will generally revert to its individual trend after positive or negative shocks. But the trend will itself be stochastic since it is the cumulative sum of the permanent residues of a lifetime of shocks. If so, and provided the shocks to different individuals are not perfectly correlated,

the health of members of a cohort will disperse over time, just as do their incomes, consumptions, and earnings.

We use data from the National Health Interview Survey (NHIS) for the 12 years 1983–94. This survey collects data on around 50,000 freshly drawn adults every year (as well as data on children), from which we use information on income, on a ordinal self-reported health status (SRHS) measure that ranges from 1 (excellent) to 5 (poor), on body-mass index (BMI), typically defined as weight in kilograms divided by the square of height in (square) meters, on race (black or white), and on sex. The arguments for using SRHS and BMI as measures of health status are discussed in the next section. Our procedure is the same as in Deaton and Paxson (1994, 1997); we create cohort data by following birth cohorts through the members of the cohorts that are randomly drawn into each year's surveys. For each year and for each cohort of men or women, black or white, we select income and health status information, from which we can create measures of central tendency, of dispersion, and of correlation for each cohort in each year. We present these data, typically as plots against age, with each cohort shown separately. We also decompose the plots into age and cohort effects, so as to isolate the trend effects that operate from one cohort to the next from age effects that are common to all cohorts as they age. A major focus is how these age patterns differ by race and by sex.

10.2 Self-Reported Health Status and the Body-Mass Index

It is difficult to define and measure a state variable that adequately captures health status during the life cycle. Much of the work on health status and income has focused on mortality, which is perhaps the only well-defined and straightforward measure, but which is useless for our purposes. Self-reported “days of illness” or “doctor visits” are themselves conditioned by socioeconomic status and often show perverse correlations with income, with better-off people apparently perceiving and treating their illnesses more seriously. Direct measures of function, activities of daily living (ADLs) and instrumental activities of daily living (IADLs), have been used to overcome these problems and provide a possible alternative to the measures used here. We start with a brief literature review that documents the links between the self-reported measures we use and other health outcomes. We conclude the section with a brief discussion of what is meant by inequality in health outcomes, and of the relevance of inequality measures derived from SRHS and BMI.

10.2.1 Self-Reported Health Status and Health Outcomes

There is a large literature that examines the relationship between SRHS and subsequent mortality. The earliest papers use data from Canada (Mossey and Shapiro 1982) and California (Kaplan and Camacho 1983). Subsequently, there have been similar studies for a variety of countries (work surveyed in Idler and Kasl 1995). Virtually all of these studies support the idea that reports

of poor health are significantly related to higher risk of mortality. Furthermore, the risk of mortality is higher for a substantial period of time (i.e., six-year and even nine-year mortality). There has been some dispute over whether self-reported health is associated with mortality in elderly populations, with the majority of researchers finding a strong association.

The standard approach in the literature has been to start by establishing the positive correlation between SRHS and mortality. Researchers then examine whether this correlation disappears on the introduction of controls for other variables, such as socioeconomic status (often quite crudely defined), as well as “objective” measures of health status and life style factors. Appels et al. summarize as follows: “Most authors are mainly concerned with the possibility that the observed associations [between self-rated health and mortality] are spurious. Eaker, for example, suggests that the female participants in the Framingham study who perceived their health as poor may have based their evaluation on their knowledge of family history of disease. Others suggest that the rating of one’s health as poor may reflect a still subclinical disease and/or an unhealthy life-style” (1996, 682). Appels et al. then go on to suggest an “alternative explanation,” that people who think of themselves as healthy build up more positive self-images, which positively affect health. The fact that SRHS is still typically correlated with mortality even after controlling for other health and lifestyle factors is often taken as support for these more psychosociological explanations, although an obvious alternative reading is that the controls are not fully effective and that people have private information about their health. For our current purposes, it is the raw correlations between self-rated health and mortality (possibly age adjusted) that are of interest, since we are trying to identify a variable that can serve as a single summary measure of health status.

The methods and results of the various studies are generally consistent, although the estimates of the size of the “effect” of poor health on mortality vary across studies, which is hardly surprising given that the groups under study often have very different characteristics. For example, in a study of people aged 70 and older, the relative risk of dying within 36 months of the initial survey is 3.5 times greater for women who report themselves in poor health relative to those in excellent/good health, and 2.5 times greater for men (Grant, Piotrowski, and Chappell 1995). Adding controls for age, education, race, marital status, ADL difficulties, and other health measures reduces the increase in relative risk to 1.5 for women in poor health and eliminates the increase in relative risk for men in poor health. A study of Lithuanian and Dutch middle-aged men indicates that, controlling for age only, those that report their health status as poor have a 23 to 80 percent increase in the risk of mortality (over 10 years) relative to those who reported good health (Appels et al. 1996).

In addition to the work on SRHS and mortality, there is growing interest in the relationship between SRHS and other health measures. Idler and Kasl (1995) find that the elderly with poor SRHS are more likely to develop ADL

difficulties. Marmot et al. (1995) use the British Whitehall II study to examine whether British civil servants who report poor health miss more work. They find that people with poor SRHS have significantly more and longer absences from work than those who report themselves in good health. Overall, the research literature supports the conclusion that SRHS is a useful health measure, in that it is correlated with and predicts health outcomes such as illness, disability, and death.

10.2.2 Body-Mass Index and Health Outcomes

The literature on BMI and health status focuses almost exclusively on the relationship with mortality. There is less consensus here than in the literature on self-reported health and mortality. Some work shows strong relationships between weight and mortality; other work shows none, or relationships only for certain groups (e.g., white women and black men, but not black women or white men; see Stevens et al. 1992a, 1992b). Many of the studies suffer from small sample sizes (which yield very few deaths) or from the use of nonrandom samples (e.g., members of a particular insurance plan or coronary heart disease study).

Our own reading of the literature is that the most convincing studies find significant effects on mortality of both very high and very low BMIs, especially for men, but that once the investigators eliminate subjects who died shortly after the survey started, or who were smokers at the time of the initial survey, the relationship between low body mass and subsequent mortality either disappears or is substantially weakened. This general result is supported by Troiano et al. (1996), which contains a literature review and “meta-analysis” of the relationship. Two specific studies that support this conclusion are Seidell et al. (1996) and Lee et al. (1993). The first studies a random sample of over 48,000 Dutch adults, aged 30–54 at the baseline, who were tracked for 12 years. (Note that because the sample is fairly young, there are still only 1,300 deaths.) The second uses a sample of over 19,000 (male) Harvard graduates for whom self-reported health information was collected between 1962 and 1966, and who were tracked until 1988. The second of these samples is obviously not representative of the overall population, but the study appears to be carefully done and uses a large sample.

The Dutch study indicates that very overweight men (defined as those with BMI greater than 30) have significantly higher rates of all-cause mortality, controlling for age. For example, those with a BMI in excess of 30 are 46 percent more likely to die than those with a BMI of between 18.5 and 25 (defined to be the baseline group). The raw data also indicate that very underweight men are more likely to die; the relative risk of mortality for those with a BMI of less than 18.5 is 2.6. However, among a sample of nonsmokers who did not die within five years of the initial period, the low-weight men do not have higher mortality. (Most of the mortality among smokers in the first five years

of the survey is from lung cancer.) The Dutch results for women are not clear-cut. Overweight women have a significantly higher risk of mortality from coronary heart disease and cardiovascular disease but do not have a significantly higher risk of all-cause mortality. A potential problem here may be that relatively few—only 500—women died, so that nothing is estimated very precisely. The Harvard study finds that, using the full sample and controlling for age only, the underweight and the overweight are at significantly greater risk of mortality than other groups. For example, with the baseline defined to be thin men with BMIs of less than 22.5, the relative risk of mortality first falls as weight rises, to 0.92 for those with BMI between 23.5 and 24.5, and then rises to 1.12 for overweight men with BMIs in excess of 26.0. When the sample is limited to those who never smoked and who did not die in the first five years, the age-adjusted relative risk of mortality increases monotonically with age and is 67 percent higher for the group with the highest BMI relative to the group with the lowest.

10.2.3 Inequality in Health Outcomes

Although the literature provides a firm basis for the relevance of SRHS and BMI as indicators of health status, it is not sufficient by itself to justify their use in the investigation of inequalities in health status. Although we may be curious to know whether BMIs become more dispersed with age—or since height varies very little, whether weight becomes more dispersed with age—we are a good deal more interested in health outcomes, so that we need to know what dispersion in BMI tells us about dispersion in health. By this token, SRHS is of more direct interest than BMI, since it contains direct information about individual welfare. Even so, there are serious difficulties in interpreting the dispersion of both measures.

Consider first a seemingly technical difficulty. When we look to see whether distributions (of income or health) are dispersing over time, the ideal criterion is that of (second-order) stochastic dominance. If distribution F_1 stochastically dominates distribution F_2 , then it will be measured as more equal by any inequality measure that satisfies the principle of transfers, effectively by any sensible inequality measure. But stochastic dominance is not preserved under monotone transformations, and our measure of SRHS is an ordinal one, so that unless we can somehow restrict allowable transforms of the 1–5 scale, we have no nonarbitrary basis for making statements about changes in inequality. The problem for BMI is less immediate but is just as serious. Because BMI is a cardinal measure, we can make well-defined statements about changes in its dispersion, or about changes in the dispersion of body weights. But since the relationship between BMI and health status is almost certainly nonlinear and possibly nonmonotonic, statements about changing dispersion in BMI have no obvious implications for changes in the dispersion of health status. One interpretation of the literature is that BMI is irrelevant for health status up

to some cutoff, say 30, after which mortality risk increases monotonically with BMI. Given such a relationship, there is no reason to suppose that statements about changes in the dispersion of BMI, well defined though they are, will have any implications for changes in mortality risk.

More serious than the problems associated with our *specific* measures is the *general* problem of what we mean when we talk about inequality in health status, and what sort of indexes might adequately capture that meaning. The possibilities are addressed by Anand and Sen (1997), with generally pessimistic conclusions. The natural (money) cardinalizations of income and consumption that permit the development of index numbers of inequality are not applicable to other concepts, such as range of functional capabilities associated with health status. Some of the literature on health inequality has focused on inequality in life expectancy—see Wilkinson (1986) for British evidence—but life expectancy, important as it is, does not capture many aspects of health status, particularly quality of life.

If we were to accept the limited goal of looking at the mean and dispersion of life expectancy, then one avenue of progress is to subject our indicators to the transformation that best predicts life expectancy, thus translating our measures into that metric. But the empirical literature, impressive as it is, is hardly adequate to establish the appropriate functional form. Even the correlations are subject to some dispute, and we are still some way from definitive conclusions about functional form (see in particular the debate on whether BMI is or is not even monotonically related to mortality). However, if we can establish that, conditional on age, life expectancy is a concave and monotone increasing function of health status as measured, and if we are concerned with average life expectancy, or with the average of any concave function of life expectancy, then an increase in the dispersion of health status is a bad thing. This is the same argument for being concerned with the distribution of income because we weight increases in income more highly the poorer the recipient. Although the point is hardly established, our reading of the literature is consistent with the view that changes in SRHS along the five-point scale have larger implications for mortality when health status is poor than when it is good or excellent. If so, increasing the dispersion of SRHS as reported here will lower average life expectancy and lower any measure based on life expectancy that values increases in life expectancy more at lower starting points. A similar argument can be constructed for BMI, at least provided we dismiss the evidence that low BMI is associated with increased mortality. For example, if we were to construct a measure such as $z = \ln 50 - \ln \text{BMI}$, life expectancy is a monotone increasing function of z and increases rapidly with increases in z at low levels—that is, among those with high BMI—becoming relatively flat thereafter. Once again, increased dispersion of z generates lower life expectancy or lower welfare if we care more about increases in life expectancy among those who have shorter life spans.

10.3 Empirical Results

10.3.1 Preliminaries

Most of our empirical results will be presented graphically, and most are straightforward transformations of the data from the 12 years of the NHIS. These surveys collect an enormous amount of information on health and on medical conditions, but very little on the economic status of individuals. In particular, there is a single family-income question, the answers to which are presented in bracketed form. The brackets are sufficiently detailed for our purposes but have the serious deficiencies that they are constant in nominal terms and that the top of the highest bracket is \$50,000, again with no change for inflation over the period 1983–94. In order to use these data, we first allocate to each individual the family income of the household in which he or she resides, using the midpoint of the bracketed range, and then convert to logarithms. Our procedures for handling the top-coding and for measuring variances are described below.

We use individuals aged 20–70, inclusive. Cohorts are typically defined by the exact year of birth, although for some of the analysis that follows we define cohorts using (nonoverlapping) five-year birth intervals, and we identify cohorts by the midpoint of their ages in a specified year. When we use exact year of age, there are 62 cohorts and 612 cohort-year cells. Not all cohorts are observed in all years because their ages must be between 19 and 71. When we use five-year age bands, there are 11 cohorts and 120 cohort-year pairs. For each cohort, sex, and, for some of the analysis, race group in each of the 12 survey years, we assemble (individual) data on the logarithm of family income, SRHS, and the logarithm of BMI. (We actually define BMI in units of pounds per inches squared rather than kilograms per meters squared, so the log of the BMI differs from the conventional measure by a constant.) From these raw data we calculate the various quantiles in the usual way on a cell-by-cell basis. In some of what follows we examine the joint distributions of health and income. To obtain means and variances of the log of income, and the covariance between income and health, taking into account the top-coding of family income, we assume bivariate normality for the pairs (log income, log BMI) or (log income, SRHS). We then fit the distributions to the data for each cohort-year cell, taking account of the censoring of log income at the log of \$50,000. This is conveniently done by fitting a Tobit model containing a constant and either SRHS or the logarithm of BMI on the right-hand side and computing moments and comoments using standard formulas for conditional normal distributions.

In some of the figures we show, not the raw data, but age effects. These are constructed by regressing the cohort-year means, variances (constructed as above), medians, or quantiles on a set of age, cohort, and year dummies. Since age is equal to the calendar year minus the year of birth, these effects must be

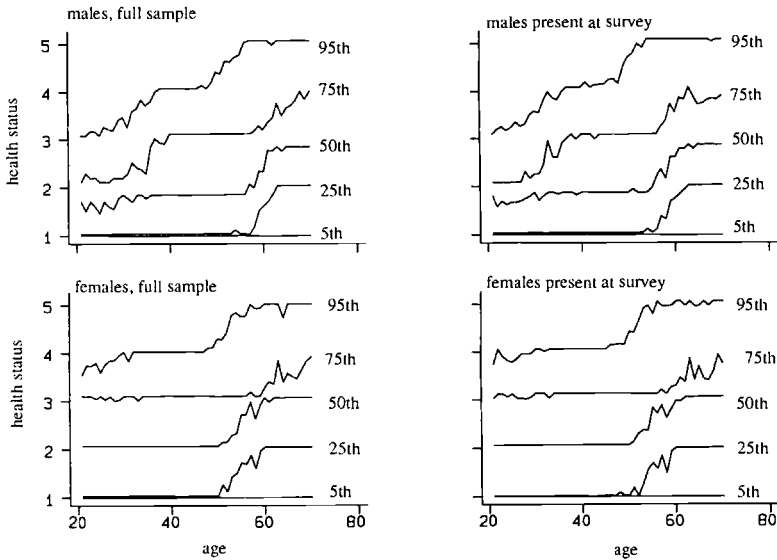


Fig. 10.1 Percentiles of health status by age (1 = excellent health, 5 = poor health)

restricted in some way. Most often this is done by omitting either year or cohort effects, and we shall explain in each case the procedures that were adopted and their influence on the results.

10.3.2 Univariate Analysis of Self-Reported Health Status, Body-Mass Index, and Income

Figures 10.1, 10.2, and 10.3 describe the univariate life cycle behavior of our three measures, separately for males and females, but pooled over all races. Figure 10.1 plots the profiles in the 5th, 25th, 50th, 75th, and 95th percentiles of SRHS, by males and by females, for all those covered in the sample and for only those present at the time of the survey. This last distinction is to allow for the possibility that reports made on behalf of others may be less reliable or systematically biased. In fact, the right- and left-hand sides of the figure are very similar, and we do not make further reference to this division of the sample. The age effects shown here were obtained by forming the percentiles for each cohort-year-sex cell and then regressing each on a set of age and year dummies; the plots show the coefficients on the age dummies. The year effects show little statistical significance, and the age effects are little affected if year dummies are replaced by cohort dummies; all of the systematic variance in these data are in the age effects, and there is little change over time at any given age.

Figure 10.1 shows that SRHS deteriorates with age—recall that 1 is “excellent” and 5 is “poor”—and that, as the initial hypothesis predicts, dispersion

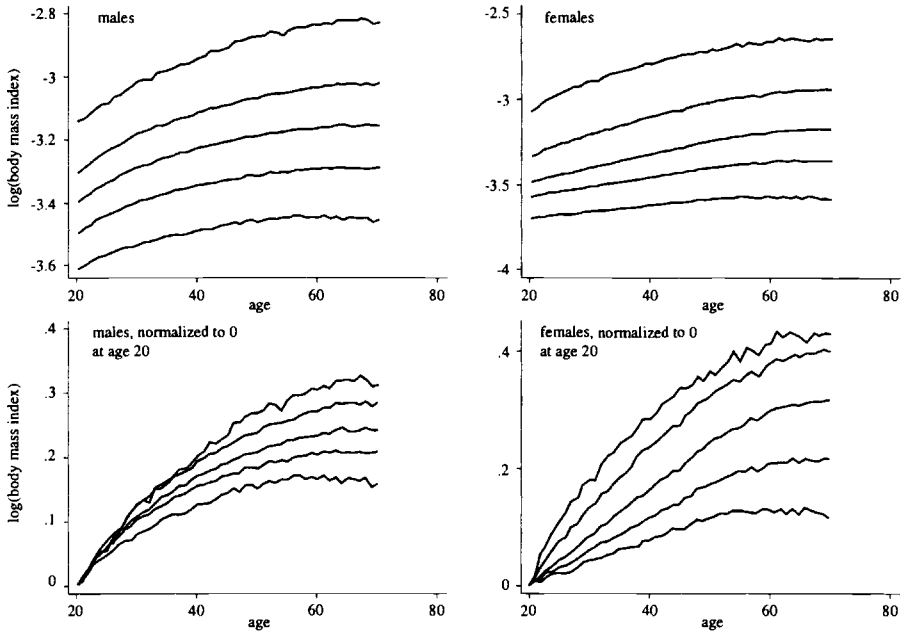


Fig. 10.2 Percentiles of log BMI by age

increases with age; see the difference between the 25th and 75th percentiles, or between the 5th and 95th percentiles. The age profiles for men and women are broadly similar, although SRHS is both worse and more variable among young women than among young men. We imagine that some of this difference is associated with pregnancy, which is not recorded in the surveys. That SRHS worsens with age is perhaps not surprising, but it implies that when people report their health status, they do not “norm” their answers with respect to the experience of those at the same age or at least they only do in part.

Figure 10.2 shows the corresponding age profiles of the same percentiles for the logarithm of BMI. The top panels—for males and females—are the coefficients of the age dummies in regressions on age and cohort dummies. There are strong cohort effects in BMI, with younger cohorts consistently heavier than their elders. Given that BMI is continuously measured, these graphs are much smoother than those for SRHS, and they also trend upward with age. At the median BMI, these graphs correspond to weight gains of about 0.3 lbs per year of age for men and 0.45 lbs per year of age for women. Women have lower BMIs than men but have greater dispersion—note the different scales on the right- and left-hand sides of the figure.

As does SRHS, BMI becomes more dispersed with age. This can be seen directly from a comparison of percentiles in the top half of the figure, but it is more clearly seen in the bottom two panels, which are constructed from the

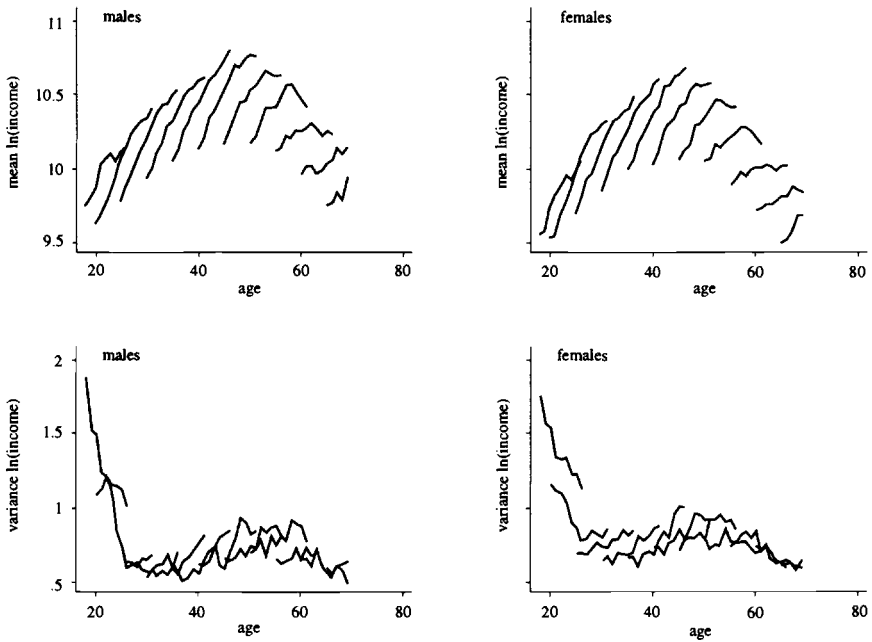


Fig. 10.3 Mean and variance of log family income by age for birth cohorts

top panels by shifting the age profiles vertically so that all are zero at age 20. The more rapid dispersion of BMI for women is then very clear in the bottom right-hand panel. But recall from section 10.2 that the links between BMI and mortality are likely much weaker for women than for men, so that their greater rate of weight dispersion may have only very limited consequences for the dispersion of health status.

Figure 10.3 plots the data on the means and variances of the logarithm of (nominal) family income, obtained from fitting the censored lognormal distributions. These figures show the raw data for each cohort, and the connected lines follow the experience of a single cohort observed year by year as it ages. The logarithm of income rises over time for each cohort and is higher for more recently born cohorts than for less recently born cohorts at the same age. The top panels also show the slower rate of growth of cohort family income in later years, a rate of growth that actually turns negative for the oldest cohorts. Note too that family incomes are lower for women than for men, a finding that does not come from distinguishing men's and women's incomes within each family—all members of a family are attributed the same family income—but rather reflects the fact that there are more women in families with lower incomes. Note finally that incomes have not been deflated for price inflation, so that the growth within and between cohorts is nominal, not necessarily real.

The bottom panels plot the estimated within-cohort variances of the loga-

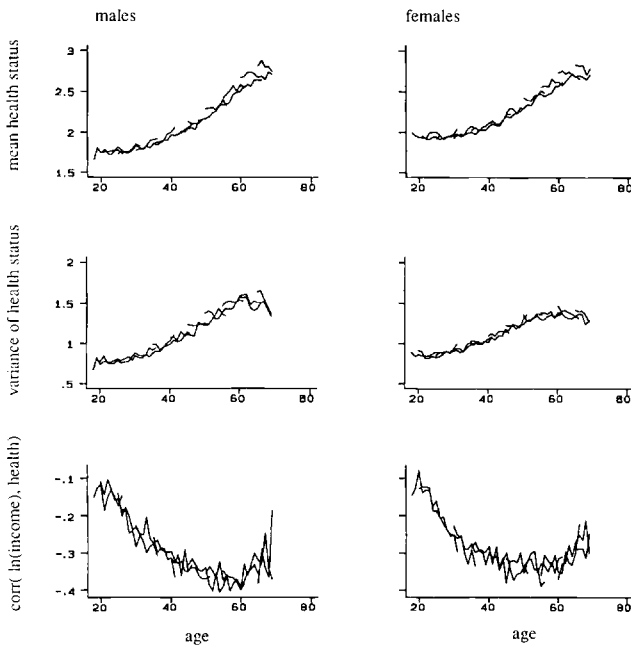


Fig. 10.4 Mean and variance of health status and correlation of income and health by age for birth cohorts

rithm of income, again on a cohort-by-cohort basis. (Unlike the levels, these variances are unaffected by inflation.) As in our previous work, which used a household rather than an individual basis, the variance of logarithms rises with age after age 25 or so until around the age of retirement, after which the variance ceases to rise or fall. The rapid falls in variance at very young ages reflect no more than the process of family formation. There are also distinct cohort effects reflecting the well-documented increases in income inequality among American families over this period.

We present this figure, less for its own interest—since it contains no information about health status—than to confirm that the income information in the NHIS, in spite of its (increasingly severe) top-coding problems and the marginality of income in the survey, can be used to reproduce the same patterns of cohort and age inequality that we obtain from higher quality income surveys, such as the Consumer Expenditure Survey and (especially) the Current Population Survey, as used in our previous work.

10.3.3 Bivariate Analysis of Health Measures and Income

Figure 10.4 shows cohort-level plots for the mean and variance of SRHS, and for its correlation coefficient with the logarithm of family income. The top two panels replicate in different form the age profiles in means and variances that we have seen in the percentile plots in figure 10.1. Most interesting here

are the two bottom plots, which document the negative correlations between SRHS and income; people with higher incomes consistently report that they are in better health. Moreover, this correlation is different at different ages; it is quite weak among those in their early 20s but becomes steadily larger (in absolute value), reaching a peak value of around -0.4 between ages 50 and 60. There are only slight differences between men and women—the correlation goes on increasing for men until age 60, whereas for women there is a plateau from around age 45 to age 60—but in both cases the correlation weakens after age 60 as SRHS deteriorates in general. This is not simply a matter of all the elderly having poor health status. As the top panels show, health status deteriorates with age, but the middle panels show only a slight decrease in the variance after age 60. It is more that, after age 60, differences in SRHS are much less well predicted by income.

These patterns of correlations between health status and income at different ages hold some clues to possible causal mechanisms. That the negative correlation should have the same age profile as the level of income (or earnings) is what would be predicted if health shocks cause income changes through participation effects or ability to work. The same health shock will have a larger effect on earnings when earnings are high, which is in the middle period of the life cycle. Against this story is the similarity of the age profiles of health-income correlation between men and women, in spite of the lower level of labor force participation among the latter.

Figure 10.5 (for males) and figure 10.6 (for females) present the correlations between SRHS and income in a way that permits us to map variance and correlation simultaneously, as well as to track different cohorts as they age. The ellipses in these figures are computed from the variance covariance matrices of log income and SRHS as follows. For each cohort-year-sex cell, we estimate the variance covariance matrix V from fitting the censored bivariate logarithmic distribution to the individual data. If z is the vector $(y, x)'$, where y is log income and x SRHS, then the points on each ellipse satisfy

$$(1) \quad z'V^{-1}z = 1.$$

The ellipse cuts the x -axis at (plus and minus) the standard deviation of SRHS and cuts the y -axis at (plus and minus) the standard deviation of the logarithm of income. The distance from the origin to the ellipse along a ray can be interpreted as the standard deviation of the corresponding linear combination of x and y . The ellipses are negatively sloped—as here—when health status and income are negatively correlated and would be positively sloped if the two were positively correlated. The ellipses can also be thought of as representing the direction and width of the joint scatter of x and y .

In both figures 10.5 and 10.6 the ellipses replicate the negative correlations in the bottom panel of figure 10.4. Each diagram shows two ellipses for the same birth cohort, one for 1983 and one for 1994, and there are diagrams for each of seven cohorts, with the youngest cohorts at the top left, and the oldest at the bottom. For the older cohorts, the ellipses are narrower and more elon-

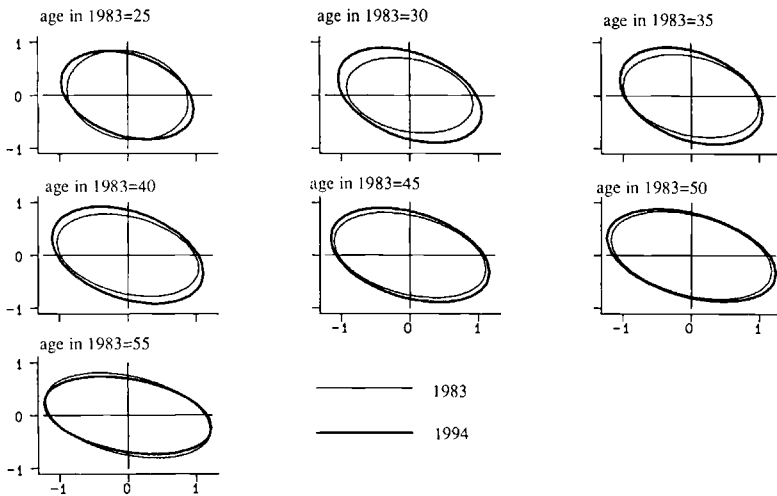


Fig. 10.5 Males, correlation of income and health by birth cohort for 1983 and 1994

Note: Log income is plotted on the y-axis, health status on the x-axis.

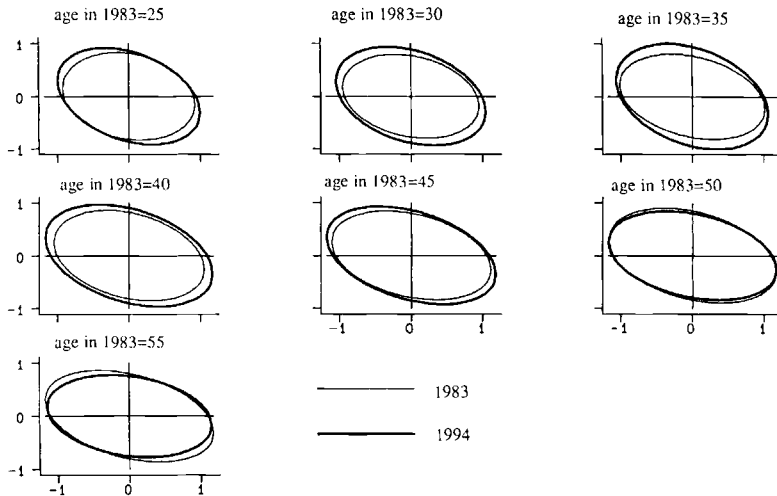


Fig. 10.6 Females, correlation of income and health by birth cohort for 1983 and 1994

Note: Log income is plotted on the y-axis, health status on the x-axis.

gated, which shows again that SRHS and income are more negatively correlated at higher ages. But within each cohort, the later (1994) ellipse—the heavier line—is typically outside the earlier (1983) ellipse as well as being more elongated. The joint distribution of income and health status becomes more negatively correlated and more (jointly) dispersed with age. This finding of

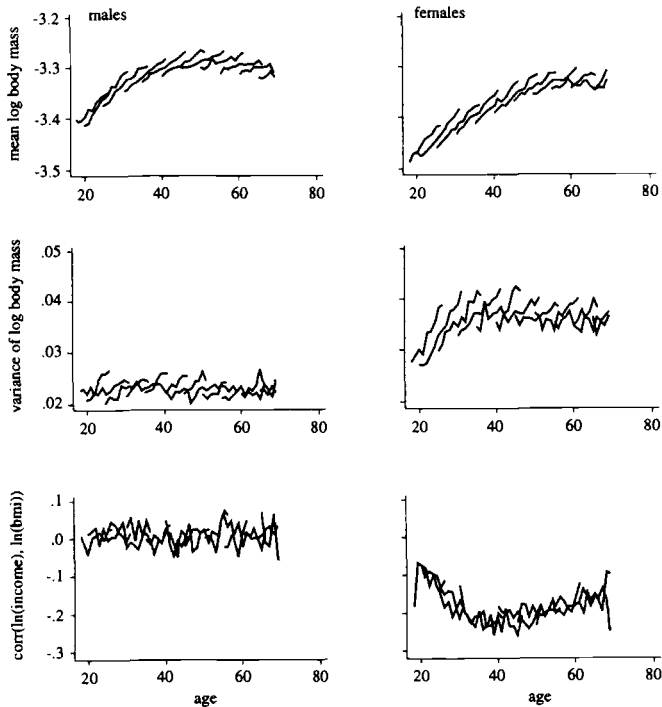


Fig. 10.7 Mean and variance of log BMI and correlation of income and body mass by age for birth cohorts

bivariate spreading generalizes and strengthens the earlier univariate findings in figures 10.1 and 10.2. Note finally that, as in these earlier findings, the rate of increase of joint dispersion diminishes with age so that, for the oldest cohorts, the earlier and later ellipses are essentially superimposed on one another.

Figures 10.7, 10.8, and 10.9 repeat the analysis with the logarithm of BMI replacing SRHS. The finding of increasing joint dispersion is replicated, but the age profiles and other patterns are otherwise quite different for the two measures. The top panels of figure 10.7 complement figure 10.2 by showing the raw cohort data for means and variance, and they repeat (with some variations) the patterns we have already seen. Weight increases with age, and there are pronounced cohort effects, with younger cohorts having higher BMIs. The age-cohort profiles of variances of the logarithm of BMI also display cohort effects; younger cohorts are not only heavier relative to height but also more variably heavy. These cohort effects obscure the positive age effects for both men and women in the middle panels, but note the generally positive slope of each cohort segment. As was the case for the interpercentile ranges in figure 10.2, the variances of logs show dispersion increasing with age, with much faster dispersion from a higher base among women.

The correlations between the logarithms of BMI and income are quite differ-

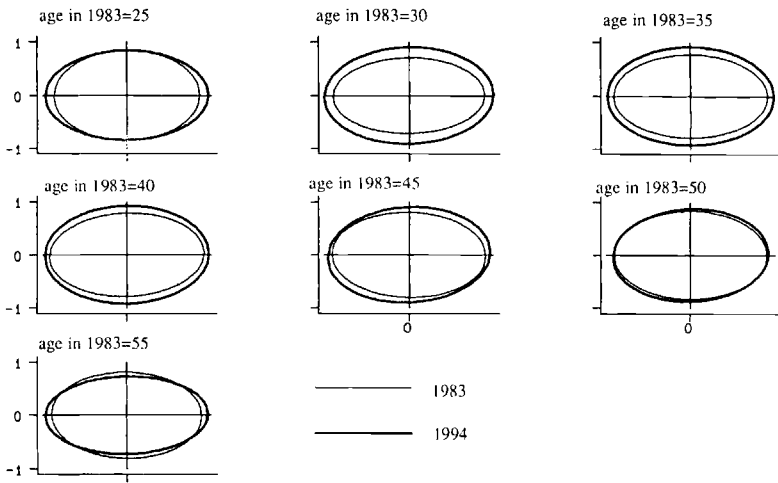


Fig. 10.8 Males, correlation of income and body mass by birth cohort for 1983 and 1994

Note: Log income is plotted on the y-axis, body mass on the x-axis.

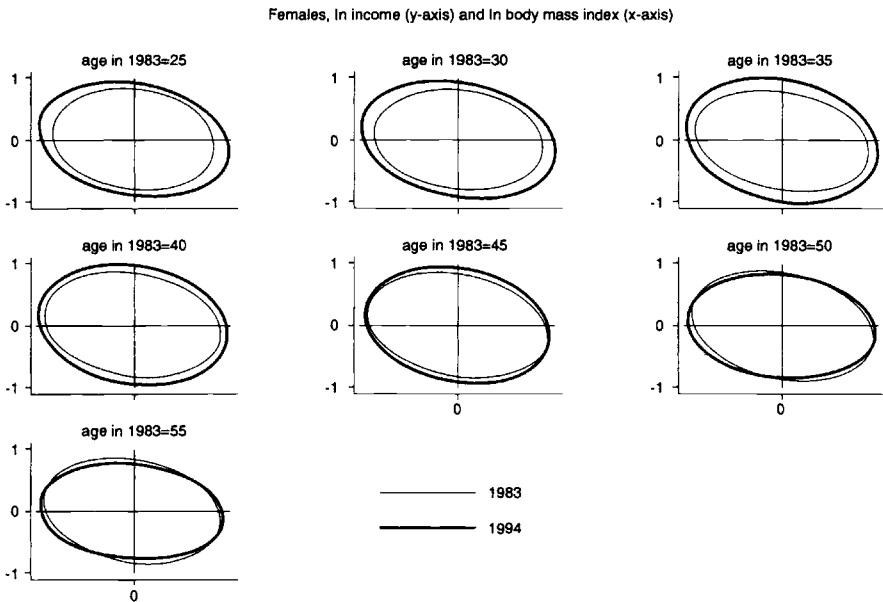


Fig. 10.9 Females, correlation of income and body mass by birth cohort for 1983 and 1994

Note: Log income is plotted on the y-axis, body mass on the x-axis.

ent from those between SRHS and log income. The bottom panels of figure 10.7 show that there is essentially no correlation between income and BMI for men at any age or for any cohort, while for women, the correlation is negative and becomes more so with age until around 40, at which age it reaches its largest (negative) value. We strongly suspect that these differences between men and women have little to do with different relationships between income and health status by sex but reflect rather the different social consequences of greater than normal weight for men and women.

In the absence of a correlation between BMI and income, the ellipses for men in figure 10.8 lie flat and, as before, move outward with age, at least among the middle-aged cohorts. Those for women in figure 10.9 are negatively inclined and show evidence of increasing joint dispersion with age among the young and middle-aged cohorts.

10.3.4 Race, Health Status, and Income

In this final subsection, we turn to differences in health status by race, and the role of income in accounting for these differences. Table 10.1 presents the raw data on SRHS by race, age, and sex. The table shows, for all years taken together, the fractions of people at each age in each of the five self-reported health categories; the numbers add to one across the rows for each sex and age. For both races, and both sexes, there is a gradual deterioration in SRHS with age. However, black males and black females are more concentrated on the right-hand side of the table than are white males and white females. At all ages, and for both sexes, there are higher fractions of whites in the “excellent” and “very good” columns, and higher fractions of blacks in the “good,” “fair,” and “poor” columns. That these differences are significant is confirmed by the very large χ^2 statistics in the final column.

The corresponding evidence for incomes is reported in table 10.2, although instead of showing fractions in each group by age, we show the fractions for five cohorts at two ages, 10 years apart, for each. The patterns are very much the same as for SRHS in table 10.1; blacks are consistently and significantly more heavily represented in the lower income groups.

The graphical analysis of these data begins in figure 10.10, which plots the age profiles of percentiles of the SRHS distribution for whites and blacks by sex. Within races, we see the same patterns as before, with (negative) levels and dispersions of health status increasing with age. But there are also differences by race, with the black distributions worse and more variable, even at early ages. Among whites aged 20–30, the median SRHS is “very good”; among the same age group of blacks, it is only “good.” Increasing dispersion—or what is close to the same thing, increasing incidence of poor health—starts at much earlier ages for blacks than for whites. A quarter of white men report themselves in excellent health until their late 50s, and a quarter of white women until their early 50s. Among blacks, the same points are reached before age 40 among males, and in the 20s for females.

Table 10.1 Fractions of People with Various Self-Reported Health Measures by Age, Race, and Sex

Age	Self-Reported Health										$\chi^2(4)$
	Excellent		Very Good		Good		Fair		Poor		
	Whites	Blacks	Whites	Blacks	Whites	Blacks	Whites	Blacks	Whites	Blacks	
Males											
20	0.522	0.424	0.293	0.245	0.158	0.272	0.023	0.053	0.005	0.007	124.9
25	0.514	0.441	0.304	0.260	0.152	0.245	0.025	0.049	0.006	0.004	86.0
30	0.491	0.398	0.305	0.281	0.160	0.248	0.036	0.059	0.007	0.014	85.8
35	0.481	0.359	0.300	0.259	0.172	0.253	0.037	0.099	0.010	0.030	173.0
40	0.427	0.321	0.306	0.276	0.198	0.287	0.052	0.076	0.016	0.040	106.7
45	0.408	0.266	0.300	0.252	0.198	0.299	0.070	0.139	0.024	0.045	132.6
50	0.367	0.199	0.282	0.229	0.235	0.341	0.077	0.146	0.039	0.085	122.6
55	0.307	0.222	0.276	0.218	0.265	0.282	0.096	0.188	0.056	0.091	80.0
60	0.263	0.156	0.250	0.182	0.281	0.288	0.116	0.229	0.090	0.146	79.3
65	0.219	0.115	0.231	0.184	0.306	0.276	0.161	0.240	0.083	0.186	73.7
70	0.192	0.105	0.225	0.172	0.318	0.266	0.179	0.301	0.085	0.157	56.5
Females											
20	0.400	0.301	0.328	0.305	0.224	0.303	0.042	0.079	0.006	0.013	112.6
25	0.407	0.325	0.338	0.306	0.212	0.285	0.037	0.074	0.007	0.010	119.9
30	0.414	0.309	0.336	0.253	0.197	0.321	0.044	0.103	0.008	0.015	286.5
35	0.403	0.275	0.323	0.261	0.212	0.308	0.049	0.125	0.012	0.030	256.5
40	0.382	0.228	0.312	0.239	0.234	0.347	0.056	0.146	0.015	0.040	341.9
45	0.360	0.186	0.293	0.242	0.251	0.313	0.073	0.197	0.023	0.062	261.8
50	0.331	0.174	0.279	0.227	0.266	0.292	0.088	0.204	0.037	0.104	246.3
55	0.265	0.127	0.258	0.187	0.312	0.343	0.118	0.239	0.046	0.104	164.5
60	0.224	0.106	0.249	0.166	0.325	0.333	0.139	0.240	0.062	0.154	142.8
65	0.185	0.086	0.242	0.168	0.337	0.299	0.168	0.292	0.067	0.154	125.0
70	0.167	0.102	0.231	0.185	0.342	0.284	0.183	0.272	0.076	0.158	71.2

Note: Chi-squares indicate tests of the equality of distributions across racial groups.

Table 10.2 Fractions of People in Different (Nominal) Family Income Categories by Age, Race, and Sex

Age	Family Income												$\chi^2(5)$
	\$0–\$9,999		\$10,000–\$19,999		\$20,000–\$29,999		\$30,000–\$39,999		\$40,000–\$49,999		\$50,000 or More		
	Whites	Blacks	Whites	Blacks	Whites	Blacks	Whites	Blacks	Whites	Blacks	Whites	Blacks	
Male													
25 in 1983	0.176	0.221	0.296	0.351	0.260	0.322	0.157	0.057	0.065	0.037	0.046	0.013	11.2
35 in 1993	0.033	0.073	0.115	0.272	0.185	0.226	0.221	0.123	0.177	0.141	0.270	0.165	24.5
35 in 1983	0.070	0.259	0.190	0.292	0.289	0.236	0.244	0.101	0.108	0.094	0.099	0.018	31.6
45 in 1993	0.047	0.219	0.081	0.203	0.126	0.138	0.165	0.097	0.136	0.039	0.445	0.305	42.0
45 in 1983	0.068	0.191	0.174	0.367	0.243	0.295	0.219	0.048	0.131	0.040	0.165	0.060	26.5
55 in 1993	0.053	0.071	0.126	0.235	0.120	0.261	0.114	0.148	0.144	0.028	0.442	0.257	17.1
55 in 1983	0.102	0.237	0.190	0.349	0.240	0.171	0.196	0.133	0.127	0.071	0.145	0.040	7.6
65 in 1993	0.051	0.238	0.207	0.489	0.226	0.117	0.185	0.117	0.090	0.000	0.241	0.039	38.4
Females													
25 in 1983	0.163	0.403	0.306	0.390	0.281	0.151	0.141	0.037	0.063	0.020	0.045	0.000	41.2
35 in 1993	0.074	0.214	0.113	0.188	0.186	0.185	0.180	0.131	0.182	0.082	0.265	0.200	44.4
35 in 1983	0.129	0.325	0.187	0.278	0.256	0.151	0.226	0.092	0.109	0.089	0.094	0.065	29.2
45 in 1993	0.041	0.216	0.108	0.200	0.127	0.169	0.127	0.101	0.154	0.088	0.443	0.225	57.7
45 in 1983	0.093	0.327	0.218	0.291	0.234	0.236	0.195	0.085	0.130	0.044	0.130	0.016	35.6
55 in 1993	0.039	0.209	0.154	0.206	0.171	0.145	0.179	0.143	0.121	0.132	0.337	0.165	44.6
55 in 1983	0.121	0.346	0.295	0.286	0.255	0.267	0.127	0.069	0.123	0.032	0.079	0.000	17.4
65 in 1993	0.110	0.221	0.266	0.417	0.206	0.139	0.166	0.123	0.103	0.012	0.150	0.089	12.3

Note: Chi-squares indicate tests of the equality of distributions across racial groups.

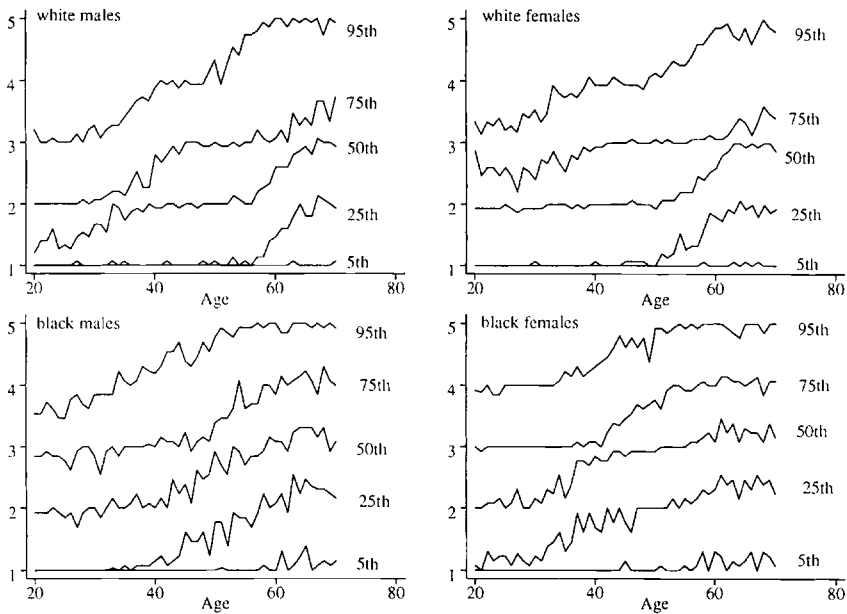


Fig. 10.10 Percentiles of health status by age for blacks and whites

Given that income and SRHS are negatively correlated, and given that blacks have lower incomes than whites, it is interesting to investigate how much of the differences in SRHS can be attributed to income, holding constant the distribution of health status conditional on income. To examine this question, we follow the analysis in DiNardo, Fortin, and Lemieux (1996) and reweight whites according to the black income distribution. The idea here is to recalculate what would have been the distribution of SRHS among whites using the actual conditional distribution of SRHS given income for whites, but with the black income distribution. Formally, if $p^w(h = i)$ is the proportion of whites whose SRHS (h) is in category i , we can write

$$(2) \quad p^w(h = i) = \sum_j p_c^w(h = i | y = j) \pi^w(y = j),$$

where $\pi^w(y = j)$ is the fraction of whites in income (y) class j and $p_c^w(h | y)$ is the distribution of health among whites conditional on income. The counterfactual that we want to create uses the white conditional distribution and the black marginal to give

$$(3) \quad \tilde{p}^w(h = i) = \sum_j p_c^w(h = i | y = j) \pi^b(y = j).$$

By comparing equations (2) and (3), we can rewrite equation (3) to give

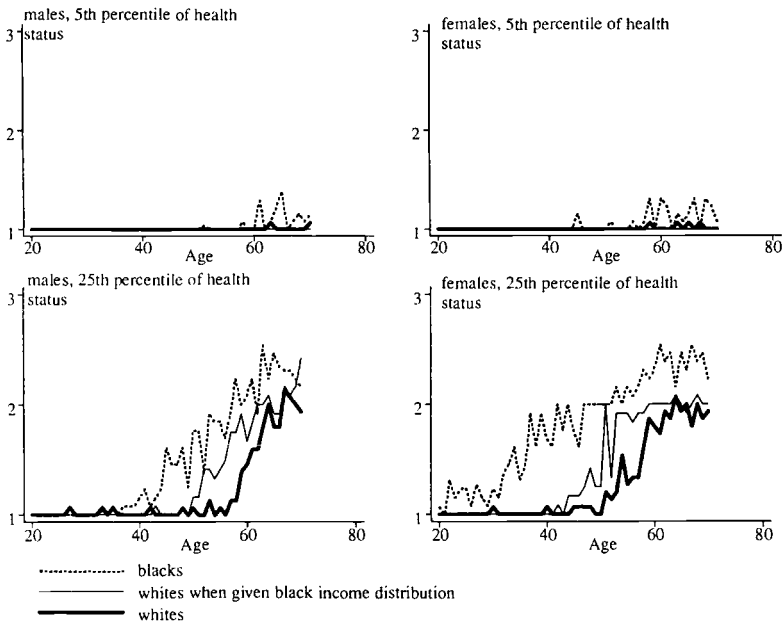


Fig. 10.11 5th and 25th percentiles of health status for blacks and whites, with and without adjusting for differences in income distributions

$$(4) \quad \tilde{p}^w(h = i) = \sum_j p_c^w(h = i | y = j) \pi^w(y = j) \frac{\pi^b(y = j)}{\pi^w(y = j)},$$

so that, finally, we have

$$(5) \quad \tilde{p}^w(h = i) = \sum_j p_c^w(h = i, y = j) \omega(j),$$

where $\omega(j)$ is a reweighting function equal to the ratio of the black to white marginal of income.

Figures 10.11, 10.12, and 10.13 show the age profiles of the 5th and 25th, 50th and 75th, and 95th percentiles of the distributions of whites, of blacks, and of whites with the counterfactual black income distribution. The general result is that income takes us a good deal of the way, but not all of the way, to explaining the difference between the two distributions of SRHS. Among those in good health—figure 10.11—the 25th percentile of the counterfactual white distribution is about half way between the 25th percentiles of the black and white distributions for men, but only a small way for women. Much the same is true for the 50th percentile in figure 10.12; a much larger fraction of the difference between blacks and white men is accounted for by income differences than is the case for women, particularly young women. At the 95th per-

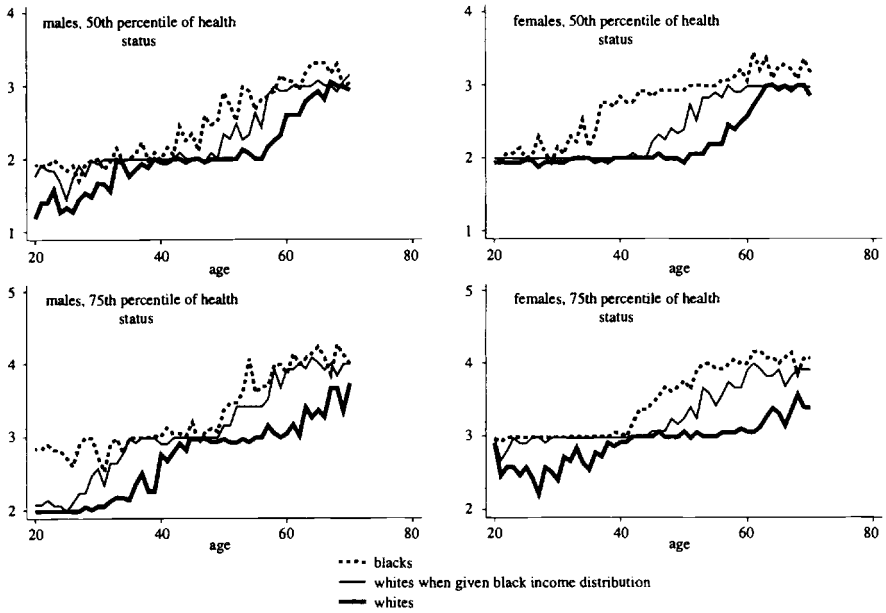


Fig. 10.12 50th and 75th percentiles of health status for blacks and whites, with and without adjusting for differences in income distributions

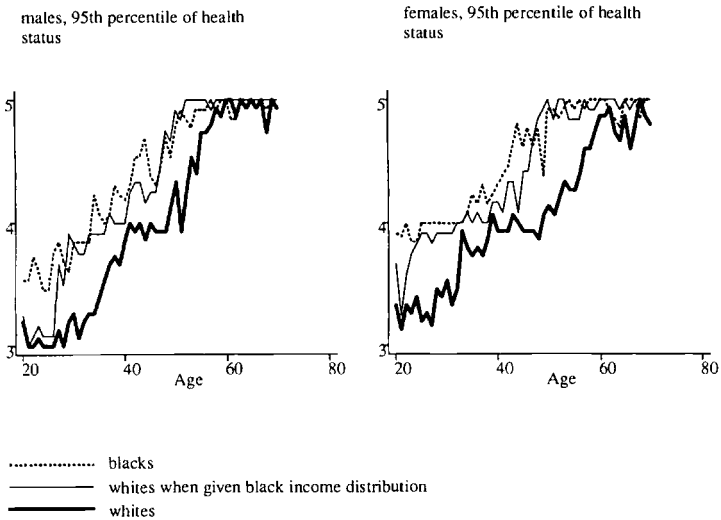


Fig. 10.13 95th percentile of health status for blacks and whites, with and without adjusting for differences in income distributions

centile, in figure 10.13, among those reporting poor health, the reweighting of the white age profile takes us most of the way to the black age profile.

The calculations for BMI are shown in figures 10.14 through 10.17. The age profiles of the percentiles of the BMI distribution are not very different between black and white men, except that the heaviest black men are a good deal heavier than the heaviest white men. In all cases, a substantial fraction of the difference vanishes when we reweight the whites to give them the black income distribution. For women, the situation is quite different. The percentiles of the black BMI distribution are at higher values of BMI at all ages for women, and only a small fraction of the difference is eliminated by conditioning on income.

10.4 Summary and Conclusions

We have presented evidence on life cycle patterns of two health-related indicators, self-reported health status and the body-mass index, as well on their relationship with income. We regard this work as exploratory; we have tried to generate stylized facts that are relevant to debates about health status, income, and inequality, even if, at this stage, there is no clear framework within which these facts should be fitted. We believe it is important to explore differences in health between people, even in the absence of an agreed methodology for thinking about inequality in health status, or even about health status itself. But by the same token, it is important to be cautious about attributing causality to any of our findings. Income and our measures of health status are linked in many different ways, through ability to pay for health, through education that is correlated with income, through lifestyle choices—such as whether to smoke and what to eat—that are conditioned by income, race, and sex.

From our findings, the following are worth highlighting:

- There is ample evidence for the proposition with which we began, that our two measures become more widely dispersed within any given birth cohort as that cohort ages. We view this as evidence in favor of a cumulative random model of health status.
- SRHS worsens with age, so that people do not report their health relative to the average health of their age group.
- The rate of dispersion with age of BMI, but not SRHS, is much more rapid for women than for men. BMI is more variable among women to start with. SRHS is more variable among young women than among young men, possibly reflecting pregnancy.
- Health status (positively measured) is positively correlated with income within cohort-year-sex cells. The correlation is lowest for the young, increases until ages 50–60, and then diminishes. BMI is uncorrelated with income for men but negatively correlated with income among women. This correlation is highest in middle age. These patterns are

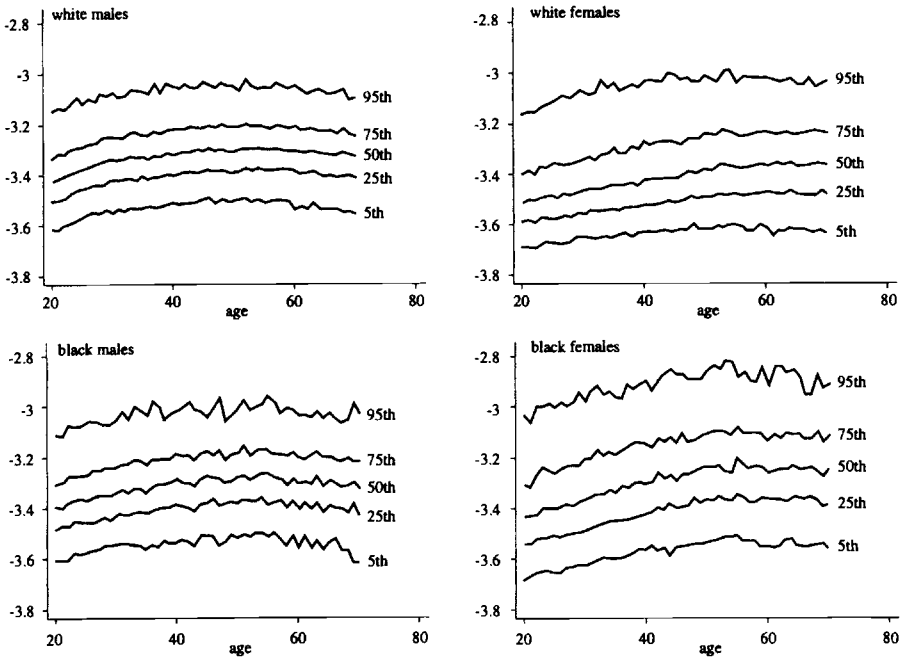


Fig. 10.14 Percentiles of log BMI by age for blacks and whites

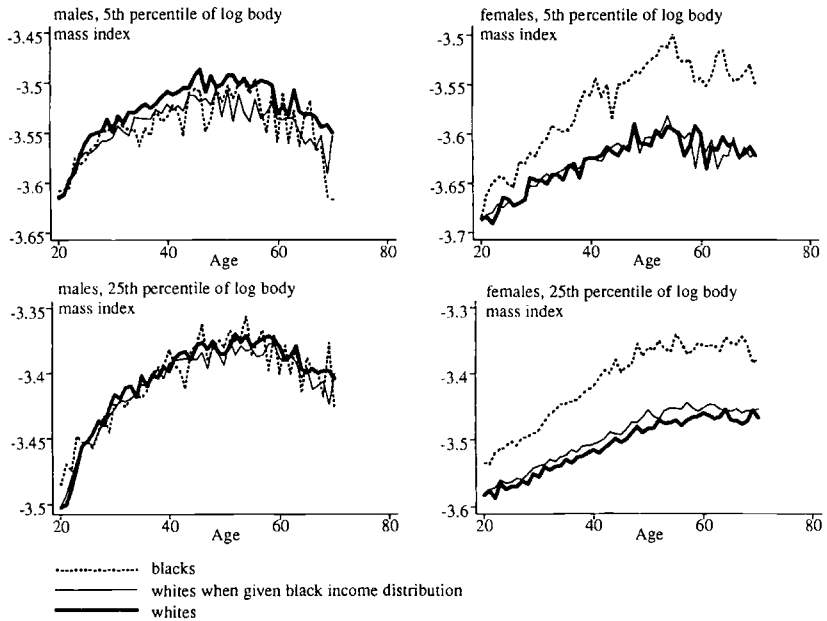


Fig. 10.15 5th and 25th percentiles of body mass for blacks and whites, with and without adjusting for differences in income distributions

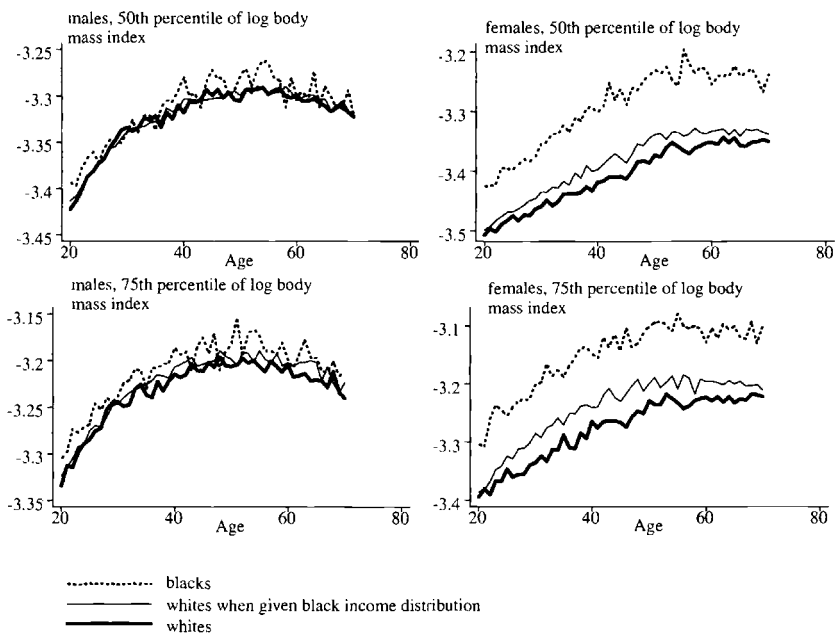


Fig. 10.16 50th and 75th percentiles of body mass for blacks and whites, with and without adjusting for differences in income distributions

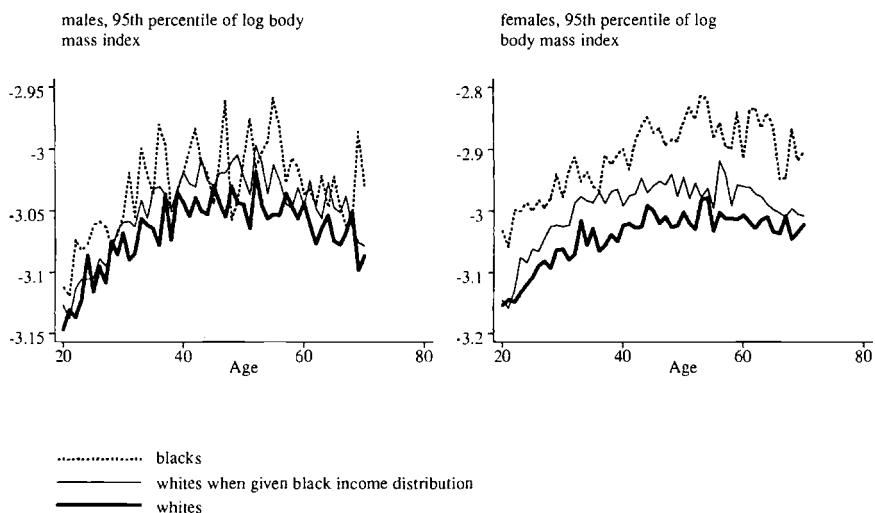


Fig. 10.17 95th percentile of body mass for blacks and whites, with and without adjusting for differences in income distributions

consistent with the hypothesis that those with lower health status earn less.

- The joint distribution of SRHS and income and the joint distribution of BMI and income “fan out” with age.
- Blacks consistently report lower health status than do whites. Some fraction—but not all—of this difference can be attributed to the lower income of blacks. Less of the difference is explained by income among women than among men, a result that is even more pronounced for BMI.

References

- Adler, Nancy E., Thomas Boyce, et al. 1994. Socioeconomic status and health: The challenge of the gradient. *American Psychologist* 49:15–24.
- Anand, Sudhir, and Amartya K. Sen. 1997. Draft report on broadening measures of inequality. Report prepared for the MacArthur Project on Inequality. Chicago: MacArthur Foundation.
- Appels, A., H. Bosma, V. Grabauskas, A. Gostautas, and F. Sturmans. 1996. Self-rated health and mortality in a Lithuanian and a Dutch population. *Social Science and Medicine* 42:681–89.
- Deaton, Angus S., and Christina H. Paxson. 1994. Intertemporal choice and inequality. *Journal of Political Economy* 102:437–67.
- . 1997. The effects of economic and population growth on national saving and inequality. *Demography* 34:97–114.
- DiNardo, John, Nicole Fortin, and Thomas Lemieux. 1996. Labor market institutions and the distribution of wages, 1973–92: A semiparametric approach. *Econometrica* 64:1001–44.
- Farber, Henry S., and Robert Gibbons. 1996. Learning and wage dynamics. *Quarterly Journal of Economics* 111:1007–47.
- Grant, M. D., Z. H. Piotrowski, and R. Chappell. 1995. Self-reported health and survival in the Longitudinal Study of Aging, 1984–1986. *Journal of Clinical Epidemiology* 48:375–87.
- Hall, Robert E. 1978. Stochastic implications of the life cycle–permanent income hypothesis: Theory and evidence. *Journal of Political Economy* 86:971–87.
- Idler, E. L., and S. V. Kasl. 1995. Self-ratings of health: Do they also predict change in functional ability? *Journal of Gerontology* 50:S344–S353.
- Kaplan, G. A., and T. Camacho. 1983. Perceived health and mortality: A 9-year follow-up of the Human Population Laboratory Cohort. *American Journal of Epidemiology* 117:292.
- Lee, I. M., J. E. Manson, C. H. Hennekens, and R. S. Paffenbarger, Jr. 1993. Body weight and mortality: A 27-year follow-up of middle-aged men. *Journal of the American Medical Association* 271:1320–21.
- Marmot, M., A. Feeney, M. Shipley, et al. 1995. Sickness absence as a measure of health status and functioning: From the UK Whitehall II study. *Journal of Epidemiology and Community Health* 49:124–30.
- Mossey, J. M., and E. Shapiro. 1982. Self-rated health: A predictor of mortality among the elderly. *American Journal of Public Health* 71:800.

- Seidell, J. C., W. M. Verschuren, E. M. van Leer, and D. Kromhout. 1996. Overweight, underweight, and mortality: A prospective study of 48,287 men and women. *Archives of Internal Medicine* 156:958–63.
- Singer, Burton, and Carol D. Ryff. 1997. Racial and ethnic inequalities in health: Environmental, psychosocial and physiological pathways. In *Intelligence, genes, and success: Scientists respond to The Bell Curve*, ed. B. Devlin, S. E. Fienberg, D. P. Resnick, and K. Roeder. New York: Springer.
- Smith, James P. 1995. Health and savings. Santa Monica, Calif.: RAND Corporation. Processed.
- Stevens, J., J. E. Keil, P. F. Rust, et al. 1992a. Body mass index and body girths as predictors of mortality in black and white men. *American Journal of Epidemiology* 135:1137–46.
- . 1992b. Body mass index and body girths as predictors of mortality in black and white women. *Archives of Internal Medicine* 152:1257–62.
- Troiano, R. P., E. A. Frongillo, Jr., J. Sobal, and D. A. Levitsky. 1996. The relationship between body weight and mortality: A quantitative analysis of combined information from existing studies. *International Journal of Obesity* 20:63–75.
- Wilkinson, Richard G., ed. 1986. *Class and health: Research and longitudinal data*. London: Tavistock.
- . 1994. Health and wealth. *Daedalus* 123 (4):61–77.

Comment David Meltzer

In earlier work, Deaton and Paxson (1994) used data from the United States, Great Britain, Taiwan, and Thailand to document that inequality in consumption, income, and earnings all increase with age. That work was of interest both because it provided a connection between demographic change and income inequality—a question that dates back at least to the work of Kuznets (1979)—and because it provided some evidence with which to test the predictions of alternative theories of resource allocation over the life cycle.

This paper extends their analysis to consider the distribution of two measures of health status: self-reported health status (SRHS) and the body-mass index (BMI). The paper begins with three justifications for examining this issue:

1. As a matter of curiosity, to investigate whether there is some general rule that dispersion increases with age—not just for income or earnings but more generally for other dimensions of human experience, for example, health status
2. Out of a particular interest in inequality in health status across age, because health is an important component of welfare (and perhaps implicitly also an indicator of welfare—income status, etc.—though this is not emphasized much in the paper
3. To a lesser extent than in the previous work, out of an interest in testing

David Meltzer is assistant professor in the Section of General Internal Medicine, Department of Economics, and Harris Graduate School of Public Policy at the University of Chicago and a research associate of the National Bureau of Economic Research.

a model of the evolution of health as a stochastic process of accumulated shocks over the life cycle analogous to the type of martingale process suggested by theories of life cycle consumption

Using data from the National Health Interview Survey (NHIS) from 1983–94, the authors decompose changes in SRHS and BMI into age and cohort effects, thereby permitting the analysis of changes in variation in these measures over the life cycle. These trends are examined by both gender and race, and their correlation with income is also examined. The paper has a tremendous number of results—and many are quite interesting—but I will focus my comments on some more general issues concerning variability in health with aging and the relationship of variability in health to welfare. I will organize my comments around the three motivations raised by the authors but will focus mainly on the first two.

I begin with the question of whether increases in variability in health status are a universal rule. It is in fact easy to determine that this cannot possibly be a general rule. Keynes knew the answer, “In the long run, we’re all dead.” Certainly, in this sense, variability in health status eventually goes to zero with age. Likewise the random walk idea, if taken literally, also falls away easily since presumably it should be two sided and we do not see people living forever or running 30-second miles at age 288. In fact we never see 90- or even 50-year-old competitive runners. Surely some components of aging are stochastic, but there are basic trajectories toward decline as well. Perhaps the relationship between age and variation in health status is best described by an inverted U.

But perhaps, even if increases in variability in health with age are not a universal rule, it is commonly the case that variability in health increases with age. In fact, this is just what one is taught in medical school. The classic example is how people respond to medications. Some older people are much like younger ones, while others are highly sensitive to medications. The possible reasons are many: changes in kidney function, liver function, lean body mass, or the sensitivity of receptors to the agent. The same lectures usually also refer to the idea of homeostasis—that the body has multiple approaches to trying to maintain equilibrium. As people age and some systems begin to deteriorate, there is at first ample reserve; then as reserves are depleted, other systems compensate. I am not sure whether biologists have articulated a model to explain increases in variability with age, but this seems to me a reasonable place to start—homeostasis and excess reserves, gradual deterioration, and then variability in observed function as weaknesses are revealed. Note that this would describe a stochastic process, but one with a powerful trajectory toward death.

There is already some evidence concerning changes in variability with age. I have not surveyed the literature concerning this; however, studies such as the Baltimore Longitudinal Study of Aging (BLSOA) (Shock et al. 1984) collected physiologic measurements—for example, of kidney function, pulmonary function, and liver function—on large numbers of people as they aged. The

BLSOA results are not consistently examined to assess variability, but it is interesting to note that one is not impressed by the increase in variability in any of these measures. Perhaps a closer look would identify further evidence of variability, and there may be questions about sample selection that may minimize variability in this study, but it would be interesting to examine these findings more fully. It would likely be valuable to review the biomedical literature to see what evidence there is for the common presumption that variability increases with age. Now I will turn to the evidence Deaton and Paxson have assembled to examine the idea that variability in health increases with age.

A first question is whether the measures of health they select—SRHS and BMI—are important measures of health. There is clear evidence that both are correlated with objective health outcomes such as mortality, and this is generally true even when a variety of other measures of health are held constant. While improvements in SRHS appear to be monotonically associated with improvements in objective health outcomes, there appears to be a U-shaped relationship between BMI and mortality, likely driven at the low end by the fact that poor health may result in low BMI. At high and even moderate BMIs, increased BMI is clearly associated with increased risk of diabetes and cardiovascular disease. The relationship at the low end of BMI aside, it seems difficult to argue that, for any individual, an improvement in SRHS or a reduction in BMI (anorexia and bulimia aside) is not generally a good thing.

Assuming that these are useful measures of health status, one question is whether the sample design of the NHIS is appropriate for looking at variability by age. Since the NHIS surveys only the living, it describes only part of the distribution of outcomes for people as they age, and of course not a random part. Likewise, the NHIS misses the institutionalized population. Presumably, the full population is more variable than the part that is responding to the NHIS. However, it seems to me that one could tell stories about the selection process into the NHIS sample that would increase or decrease variability in this sample compared to the full population over time. It is probably most likely that variability increases more over time than suggested by the NHIS due to selection out of the sample by death and by institutionalization. Another point worth remembering is that the relationships between income and health status may reflect causation in either direction. The effects may be largest at the middle ages and smaller at young and old ages because it is in the middle ages that a shock to health has the biggest effect on earnings. This could also explain why much more of the differences in SRHS between black and white men are accounted for by income. Likewise, it could explain why less of the differences in SRHS between black and white woman are explained by income, since women are less likely to be primary earners. Though this is suggestive, there is a long way to go to pin down such interpretations. Studies of specific diseases may offer creative ways to disentangle these two effects.

Since this is an exploratory paper I want to take advantage of the opportunity to throw out some ideas about expectations concerning health and health in-

equality more generally and try to make a further case for why we should all think hard about inequality. After this I want to come back to the central topic of the paper for a bit and think about SRHS again in the context of some of the ideas I will raise.

Let me begin with a postulate: that in health and medical care it is not only absolute levels but levels relative to some set of expectations that matter to people's welfare. The idea that expectations matter is surely not a new one; sociologists have argued this for years, and even economists such as Richard Easterlin have argued for the importance of such expectations in generational conceptions of welfare. But I think the case is ultimately more compelling in health care. In health care, the fact that life is at stake makes both the formation of expectations and the psychic penalties associated with deviation from expectations all the more salient. There is a powerful sense that we should do everything possible to preserve life, even at times when the quality of that life may be quite poor. The Bible tells us we should not kill. Perhaps this is why the doctor bringing news of a patient's death to the family reassures them that "we did all we could." There is no economic model to explain the solace in that remark other than a psychic return to knowing that an expectation about how life is to be valued was met—the patient is surely no less dead. The importance of expectations is also revealed by a similar phenomenon, that of "laying crepe" with the family of a sick patient: by preparing them for the possibility of their loved one's death, it is somehow made more tolerable.

I have been thinking recently about the implications of such a model of the role of expectations in health and in health care and think they are worth mentioning before discussing how they apply to this paper. The two most interesting implications relate to the welfare consequences of technological change and the policy implications of social determinants of expectations. It is in the latter that ideas of social determinants of health are perhaps most salient.

Let me start with the technical change idea. Though we are concerned about the costs associated with new technologies, the general presumption suggested by standard economic models is that they must be welfare enhancing. Presumably, free disposal ensures this. However, allowing expectations into the model changes this. Consider a model where expectations (H) shape people's utility from health (h), so that the goal is to minimize the gap (G) between expectations and health: $U(X, G) = U(X, H - h)$. What does technical change do in this case? The answer depends on the nature of the technical change—whether it is frontier enhancing or cost reducing. If it is frontier enhancing and this produces heightened expectations, it reduces welfare. In this case people will spend more yet be less satisfied. This is what Arthur Barsky has called the "paradox of health"—increasing dissatisfaction with health in the context of increasing capabilities (Barsky 1988). Only if technological change is cost reducing will it unambiguously increase welfare.

This result assumes that expectations are driven by technical change—that is, by the maximum technologically feasible. But let me begin to develop the

link to this paper. Expectations are likely to be driven less by the extremes possible in theory than by the experiences one observes. If you have a cataract and your friends all had their cataracts removed, you expect to have yours removed. If people over age 55 do not get in vitro fertilization (IVF), you may not expect to get IVF if you are above that age. An interesting example of this is with respect to eligibility for organ transplantation. Often the rules are written more strictly than they are applied in practice. One possible reason is to lower expectations.

But if expectations are formed based on the actions and experiences of others, what does this mean for health care economics and policy? One key implication is that my personal decision about my health care affects your welfare. If I get a transplant—even by paying for it myself—you want one when you get sick. If I live to 90, you want to live to 90. As an older but still quite young friend of mine with a serious illness said recently: “It does not bother me so much that I am not completely healthy—but that it is at an age when all my peers are still so healthy.” In the extreme, variability is the sole measure of aggregate welfare.

The economic implication is that with goods whose value is determined in such a social context, there is an externality associated with consumption. This implies that markets may not be efficient and that private interest may not serve the social good. I do not know if this is idle speculation or a truly important aspect of how we as humans perceive our welfare with respect to health. That seems to me an empirical question and one I do not know quite how to test. Perhaps studies such as this that examine inequality are a good place to start. But if it is true that expectations concerning health and health care are indeed important and socially determined, the implications are surely profound.

If we take seriously this idea of expectations and think about the meaning of SRHS, there are, in fact, multiple concepts of SRHS, differing in how they address the role of norms, that are important to consider. In the NHIS, respondents are simply asked to rate their health without specific reference to a comparison group; but in many other studies, respondents are explicitly asked to compare themselves to others of their age. In a recent review, Ellen Idler and Yael Benyamini conclude that it does not matter much whether people are *told* to make the comparison with people their own age: “It is possible that the comparisons with socially similar others are implicit in the cognitive process that produces these ratings; if so, directing the respondent’s attention in this way would be redundant, which it appears to be” (1997, 30). A quote from a respondent helps illustrate: “My leg. That’s the only thing that’s holding me back. I feel good. And when I look around . . . I’m not sick. Believe me, some of these men and ladies around here . . . I’m not sick. I don’t wanna brag, but I wouldn’t wanta be the way some of these people here are that hafta be here.” How a person defines the group in reference to which norms are formed is clearly complicated. This is illustrated by the response of one 85-year-old woman asked to compare herself to others her age:

Interviewer. Is it hard for you to compare your own health with that of other people your own age?

Respondent. Well most of them are dead, aren't they?

Thus it seems likely that SRHS may be age normed; however, we do not know for sure. Assume that SRHS is not at all normed by age. Then changes in SRHS reflect changes in health. In that case I begin to worry about floor and ceiling effects. Maybe young people vary a lot in health, but perhaps it is just in degrees of excellent health. This is an inherent problem with an ordinal and closed-end scale. On the other hand, if SRHS *is* age normed—and the cognitive psychologists studying it seem to think so—then it is not clear what variability means. The paper talks about the fact that second-order stochastic dominance—and therefore the meaning of standard measurements of variability—is not preserved under monotonic transformations, and in fact the evidence for SRHS is that we have exactly that sort of norming. This is the attraction of physiologic measurements, ranging from BMI to lung and kidney function, and more objective measures of functional status such as activities of daily living. I hope this paper will push us to probe more deeply into those issues.

References

- Barsky, A. J. 1988. The paradox of health. *New England Journal of Medicine* 318 (7): 414–18.
- Deaton, Angus, and Christina Paxson. 1994. Intertemporal choice and inequality. *Journal of Political Economy* 102:437–67.
- Idler, Ellen, and Yael Benyamini. 1997. Self-rated health and mortality: A review of twenty-seven community studies. *Journal of Health and Social Behavior* 38:21–37.
- Kuznets, Simon. 1979. *Growth, population, and income distribution: Selected essays*, 1st ed. New York: Norton.
- Shock, Nathan, et al. 1984. Normal human aging: The Baltimore Longitudinal Study of Aging. NIH Publication no. 84-2450. Washington, D.C.: Government Printing Office, November.