3 Health Improvement and Income Growth in the Long Run

David N. Weil

The rise in the material standard of living that has taken place over the last 250 years is rightly celebrated as an enormous accomplishment of our species. But parallel to this increase in the amount of goods and services that people consume has been an improvement in human health, both the number of years lived and health status during those years. In terms of welfare, it is not immediately clear whether income or lifespan growth is the greater accomplishment.

In this essay I explore the relationship between economic growth as conventionally measured (in terms of income or consumption) and improvements in health. I begin with a discussion of how income and health have evolved over time in the leading countries (that is, the richest and healthiest). This is followed by an examination of the cross-sectional relationship between health and income among countries. I then move on to discuss the causal relations between health and income, both over time and across countries. Among other things, I argue that the very strong cross country relationship between health and income may reflect relatively little causation running in either direction, but is instead strong evidence of the importance of "institutions" for both outcomes.

3.1 Basic Facts

3.1.1 Long-Run Changes in Income and Health

Economists studying economic growth are well acquainted with the basic facts regarding the long-run evolution of the standard of living. First, for most of human history, there was almost no long-run trend in output per capita, as best we can measure it. Standards of living in the most advanced countries on the eve of the Industrial Revolution might have been a bit higher than they were in, say, Roman times, but the difference was small, and on an annualized basis, the rate of growth was trivial. Keynes' (1930) characterization was largely correct: "From the earliest times of which we have record—back, say to two thousand years before Christ—down to the beginning of the eighteenth century, there was no very great change in the standard of life of the average man living in civilized centres of the earth. Ups and down certainly. Visitations of plague, famine, and war. Golden intervals. But no progressive, violent change. Some periods perhaps 50 per cent better than others—at the utmost 100 per cent better—in the four thousand years which ended (say) in A. D. 1700." Such relative constancy in living standards is well explained by a combination of slow technological progress and the Malthusian channel by which temporary increases in income called forth higher population growth, which in turn pushed living standards back toward their equilibrium level (Galor and Weil, 2000). A second well-established fact is that, among the countries that first entered into the phase of modern economic growth, the trend growth rate of income has been remarkably constant for approximately the past 150 years. In other words, plotted on a log scale, income per capita looks to be very close to a straight line. Of course, there are important deviations from this trend, such as the Great Depression. Further, careful examination reveals evidence that the trend in growth is not completely constant. Finally, it should be noted that between the period of near zero growth that characterized most of human history and the current regime of constant growth there was a period of a century or more during which the engine of growth was revving up but had not yet reached full speed.

While the growth of the leaders has been remarkably constant, a great deal of action has occurred elsewhere in the cross-country income distribution. Many countries experienced little or no growth during this 150-year period, creating enormous gaps between the richest and poorest countries. Other countries entered late into the growth game, and their growth has followed a different pattern than that of the leaders. As Lucas (2000) points out, the later a country began the process of growth, the less subject it was to the speed limit that restrained the leaders. Growth miracles have always been of the catch-up variety, with growth inevitably slowing down as trailing countries approached the front of the pack. Also, while growth in the leading countries has been relatively constant, a good number of growth disasters and stagnations affected countries that lagged behind.

The basic facts regarding health improvement happen to be startlingly similar to those just described for economic growth. First, as with income, there was apparently little long-run change in average health during most of human history. Parallel to the previous quotation from Keynes, we can in this case quote Malthus (1798): "With regard to the duration of human life, there does not appear to have existed from the earliest ages of the world to the present moment the smallest permanent symptom or indication of increasing prolongation." As with income growth, the exact timing of the switch from stagnation to growth is disputed. At one extreme, de la Croix and Licandro (2015) use biographies of three-hundred thousand famous individuals born starting in the twenty-fourth century BCE as a tool to measure the trend in mortality. By construction, their data focus only on adults who lived long enough to become prominent and on regions that were sufficiently developed such that written biographical records survive, primarily Europe. De la Croix and Licandro date the beginning of mortality improvements to the cohort born between 1640 and 1649. The mean lifespan of famous people was sixty years in the four millennia prior to that cohort, and by the time of the cohort born in 1869, the mean lifespan had risen to sixty-nine. By contrast, Peltzman (2009) characterizes the history of life expectancy in five currently rich countries as showing little progress until approximately 1850. Maddison (2005) reports life expectancy in Roman Egypt at twenty-four years, which is the same as the value for England during the fourteenth century. By the middle of the eighteenth century, life expectancy was thirty-five years in England but still twenty-five years in France.

As with income growth, a second important fact regarding health improvement is the relative constancy of health improvement in the cutting-edge countries. Oeppen and Vaupel (2002) find that life expectancy in the countries with the highest life expectancy in the world has increased linearly since 1840 at a pace of three months per annum, with no sign of a slowdown (although there has been more fluidity in the identity of the leading countries with respect to life expectancy than there has been in the case of income). In addition, similarly to income, the constant pace of improvement at the cutting edge for life expectancy has contrasted with more variability among countries with lower life expectancy. The initial increase in life expectancy occurred only in a few countries (generally the same ones that experienced early increases in income), resulting in historically unprecedented gaps in health between countries. Also, countries with life expectancy well below that of the leaders have experienced health miracles that were, if anything, more dramatic than the growth miracles with which economists are familiar. A prominent example is Sri Lanka, where life expectancy rose by twelve years during the period from 1946 to 1953. During a longer period, 1950 to 1999, life expectancy rose by twenty-eight years in India and by thirty years in China, which in both cases was more than twice as fast as the pace in the cutting-edge countries.

Although there are many interesting aspects of the income and health data just described, this essay focuses on the first fact previously listed for each of these dimensions—the relative constancy of growth at the leading edge for approximately the past 150 years. Before beginning, however, it is worth taking a moment to note how unexpected and even strange the constancy of these two growth rates is. Regarding income, the odd constancy of growth is best highlighted by considering all the aspects that influence the growth rate of income in the standard models of growth. The brief list of such factors is technological progress, accumulation of factors of production (physical and human capital), and change over time in the institutional framework in which production takes place. In some of these cases, we have well-established quantitative frameworks for thinking about how a given change will contribute to growth. For example, within the context of the Solow model, we can trace out the growth trajectory that will result from a given increase in the investment rate or decrease in the population growth rate. Similarly, the Mincer framework gives a quantitative answer to the question of how much a particular observed rise in educational attainment should affect growth. Regarding technological progress, we have more theoretical models of the function that maps R&D effort into technological progress and that study the conditions under which the growth rate of technology will be constant. By contrast, concerning institutions, we we do not have a workable quantitative framework for either measuring or estimating the contribution of institutional change to growth. What is clear, however, is that in considering the past 150 years of growth in the leading countries, all of these mechanisms should have been operative. Population growth has fallen; investment has risen; educational attainment has risen massively; investment in new technology, measured by the fraction of the labor force doing R&D, has increased; and institutions have changed beyond recognition (consider, e.g., the share of output controlled by the government). Further, the changes in these various growth determinants were far from constant during the period under consideration. Investment

rates, for example, rose considerably more during the first half than during the second, while the rise in R&D spending is primarily a post–World War II phenomenon. Why these diverse changes resulted in relatively smooth growth is basically a mystery (Jones, 2002). Furthermore, in the face of this mysterious outcome, it is hard to know how confident one should be in extending the trend into the future.

Once again, the case of health improvement is largely parallel. Rather than looking to a production function, in this case, we see changes over time in both ages and causes of death as well as the means by which health improvements were achieved. In the United States, 64% of the gain in life expectancy that occurred between 1900 and 1940 was attributable to reductions in mortality before age fifteen, while 7% was due to reduced mortality between ages forty-five and sixty-four and only 3% to reductions in mortality at ages sixty-five and greater. Between 1960 and 1990, reduction in morality in the first age group contributed 30% to increased life expectancy, compared to a 28% contribution from the second group and 36% from the third. Going forward, the share of life expectancy gains resulting from decreased mortality among the elderly will continue to increase.[1] In 1900, the leading causes of death in the United States were pneumonia and influenza (together, 11.8% of deaths), tuberculosis (11.2%), and diarrhea and other intestinal conditions (8.3%). In 1998, infectious diseases such as these had largely been eliminated, and the three top causes of death were heart disease (31.0%), cancer (23.1%), and stroke (6.8%).[2] In the leading countries, improvements in life expectancy before the late nineteenth century are considered largely the result of rising living standards, although there is some debate about this issue (Fogel 1993; Deaton 2013). In the last decades of the nineteenth century and into the first several decades of the twentieth, the main drivers of rising life expectancy were improved access to clean water and other changes in behavior that followed the acceptance of the germ theory of disease. From the mid-twentieth century onward, explicit medical interventions rose to a dominant position as a source of further gains in longevity.

Given the relative constancy of income growth and life expectancy improvement in the leading-edge countries, it will be useful to form a simple ratio between the two. As noted previously, life expectancy has increased by roughly three months per year at the cutting edge; for income growth, a good rough-and-ready figure is 2% per year. (This is the figure used by Lucas [2000]. The value for the United States might be closer to 1.8%.) For comparability with the following analysis, I convert these numbers into a ratio of life years per one log unit of GDP. Since a country with GDP per capita growing at 2% per year will increase the log of GDP per capita by one in fifty years, we simply multiply the quarter-year gain in life expectancy per year by fifty. Thus, the ratio is 12.5 life years per one log unit of GDP.[3]

Figure 3.1 presents an alternative reckoning of similar data. I show the log of income per capita and life expectancy at the decadal frequency from 1960 to 2010 for eight countries that should be at or near the cutting edge in terms of both income and health. Several points are notable. First, if we ignore the first decade of data (1960–1970) in a few of the countries that were arguably still experiencing catch-up growth (Japan, Italy, and France), then the facts mentioned earlier regarding income growth and life expectancy improvement being roughly constant seem to hold reasonably well (with the exception of growth being unusually slow from 2000–2010). Second, it is interesting to note that there are persistent level differences among the income-life expectancy curves that are traced out: the curve for the United States is persistently lower than that for Canada, which is in turn lower than that for Japan. The slope of the income-life expectancy curve is not exactly the same as that based on the calculation from Oeppen and Vaupel described

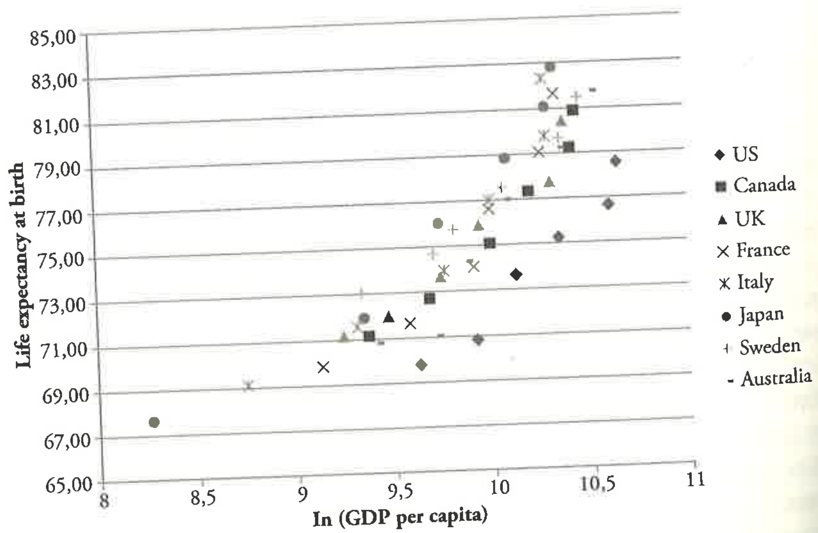

Figure 3.1
Income and life expectancy, 1960–2000.

previously, but it is not very different either. In this data, the ratio is about nine years life expectancy gain per one log unit of GDP.

3.1.2 The Cross-Sectional Relationship between Income and Life Expectancy

A good deal of the academic discussion of the causal relationship between health and income focuses on the cross-sectional relationship between the two. Preston (1975) was the first to examine this relationship in detail, and it is still called the Preston Curve. In his original article, Preston notes that the relationship between life expectancy and income has been changing over time—with life expectancy plotted on the vertical axis, the Preston Curve has been shifting upward. The increase in life expectancy experienced in a country over time can be decomposed into the part due to higher income (movement along the Preston Curve) and shifts in the curve itself. Preston's calculation was that less than one-third of the average mortality improvement observed between 1930 and 1960 was due to movement along the curve, with the remainder due to shifts in the curve. In other words, income gains were not the primary source of health improvements.

Figures 3.2 and 3.3 present data on life expectancy and GDP per capita for 1960 and 2012. Data are from the World Development Indicators (WDI) database. Points in the scatterplot are scaled proportionally

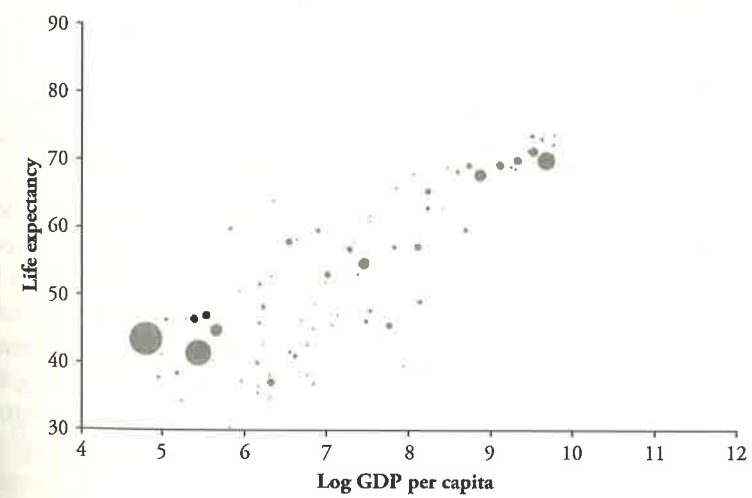

Figure 3.2
Preston Curve for 1960.

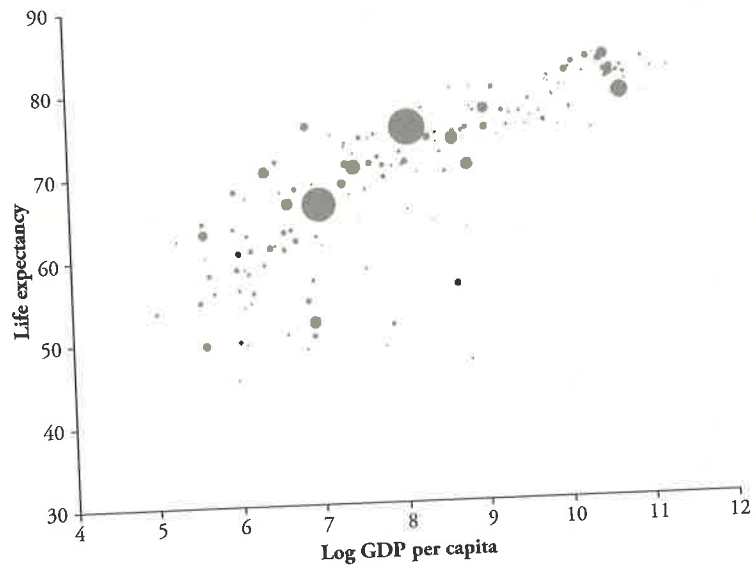

Figure 3.3
Preston Curve for 2012.

to the population of the countries they represent. Comparing the pictures, the upward movement of the Preston curve is clear. A log-linear specification seems to fit the data relatively well.

One would like to ask of this data both how much the Preston curve has shifted upward and whether the slope of the Preston Curve has changed over time. A natural approach to this question, especially given the way in which the figure is drawn, would be to regress life expectancy on the log of GDP. A moment's consideration, however, raises the question of why one would want to run the regression this way rather than the reverse direction, that is, regressing log GDP per capita on life expectancy. Since both variables are clearly endogenous and no claim is being made that the Preston Curve is somehow a structural relationship, there is not a good case for doing it one way in preference to the other. Instead, I calculate the slope of the curve by running the regression both ways, inverting the estimate from the second regression to put it in units that are similar to the first, and then taking the geometric mean of the two.

A final issue to address is weighting. Many of the observations in the World Development Indicators (WDI) database represent very small

countries, such as Bermuda (population sixty-five thousand). If one views every country as a draw from some distribution, and if measurements in small countries are just as accurate as those from large ones (which would probably be the case if were to compare Bermuda and China), then there is no statistical reason why observations from different countries should be treated differently based on population. Conversely, to find an empirical relation that describes most of the people in the world, it is natural to weight large countries more heavily. Following, I take both approaches, presenting results for unweighted regressions as well as results using population weights (aweights in STATA).

Table 3.1 shows the results. Panel A presents regressions of life expectancy on log GDP per capita, while Panel B presents regression of log GDP per capita on life expectancy. In Panel C, I show both the inverted coefficient from the Panel B regression and the geometric mean of the slope implied by Panels A and B.

Using the unweighted data, the Preston Curve became appreciably flatter between 1960 and 2012. Using the geometric mean of the estimates, the reduction is 2.7-years life expectancy per log unit of GDP per capita. This flattening is seen in whichever direction the regression is run. Using the weighted data, the evidence for a decline in the slope of the Preston Curve is weaker. Using the geometric mean, the decline is only 0.8 years per log unit of GDP per capita, and regressions run different ways actually give different signs for the direction of the change. In both the unweighted and weighted regressions, the R-squared falls between 1960 and 2012, although the change is more pronounced in the weighted regressions.

For 2012, the unweighted and weighted estimates of the slope of the Preston Curve are very similar: 5.88 and 5.59 life years per log unit of GDP per capita, respectively. For 1960, the estimates differ appreciably: 8.59 and 6.41 life years per log unit of GDP per capita (these slopes refer to the geometric mean described previously). By any of these measures, the slope of the Preston Curve is flatter than the slope of the income-life expectancy trajectory traced out by the data for the leading countries, which as discussed earlier, has a long-run slope of 12.5. The reason for this difference is, of course, that the Preston Curve shifts upward over time.

The facts presented previously can be combined in a single picture. Figures 3.4 and 3.5 show the world roughly as we have experienced it, with the richest countries experiencing growth in income and life expectancy at rates of 2% per year and one-quarter year per year,

Table 3.1
Preston Curves for 1960 and 2012

(A) Regression of Life Expectancy on Log of GDP per Capita

	1960		2012	
	Unweighted	Weighted	Unweighted	Weighted
Ln(GDP/capita)	7.15 (0.51)	5.94 (0.26)	4.59 (0.27)	4.42 (0.26)
Obs.	88	88	181	181
R-squared	0.69	0.86	0.61	0.63

(B) Regression of Log GDP per Capita on Life Expectancy

	1960		2012	
	Unweighted	Weighted	Unweighted	Weighted
Life Expectancy	0.0968 (0.0070)	0.1447 (0.0063)	0.1327 (0.0080)	0.1415 (0.0082)
Obs.	88	88	181	181
R-squared	0.69	0.86	0.61	0.63

(C) Derivation of Preston Curves: Life Expectancy per unit Log GDP per Capita

	1960		2012	
	Unweighted	Weighted	Unweighted	Weighted
From Panel A	7.15	5.94	4.59	4.42
From Panel B	10.33	6.91	7.54	7.07
Geometric Mean	8.59	6.41	5.88	5.59

respectively, and the Preston Curve changing in slope according to the estimates presented and shifting upward. I stress that this is a stylized representation (along the lines of Lucas 2000). Specifically, I take as the anchor for the 2012 Preston Curve the values for Japan: log GDP per capita of 10.33 and life expectancy of 82.8 (the highest in the world among major countries). I draw the Preston Curve for 2012 under the assumption that it goes through this point and has a slope equal to the geometric mean of the estimates presented previously. To construct the Preston Curve for 1960, I assume that, in Japan, life expectancy and GDP per capita grew at the stylized rates mentioned earlier, specifically, one-quarter year per year and 2% per year. These were not the actual

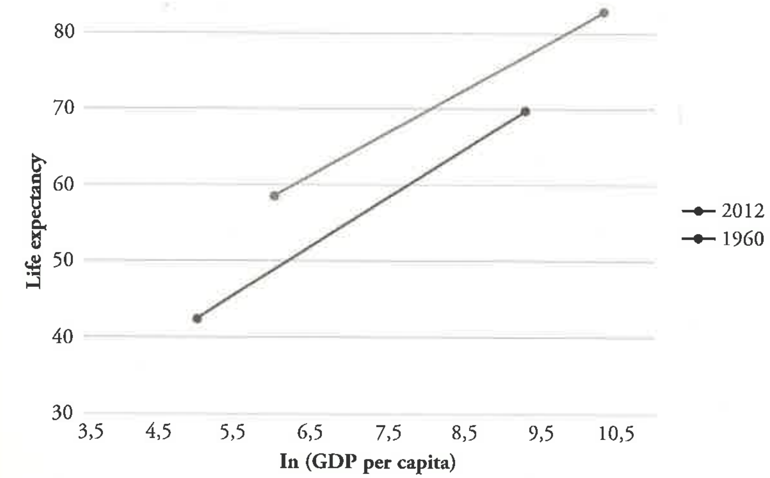

Figure 3.4
Stylized Preston Curves (weighted estimates).

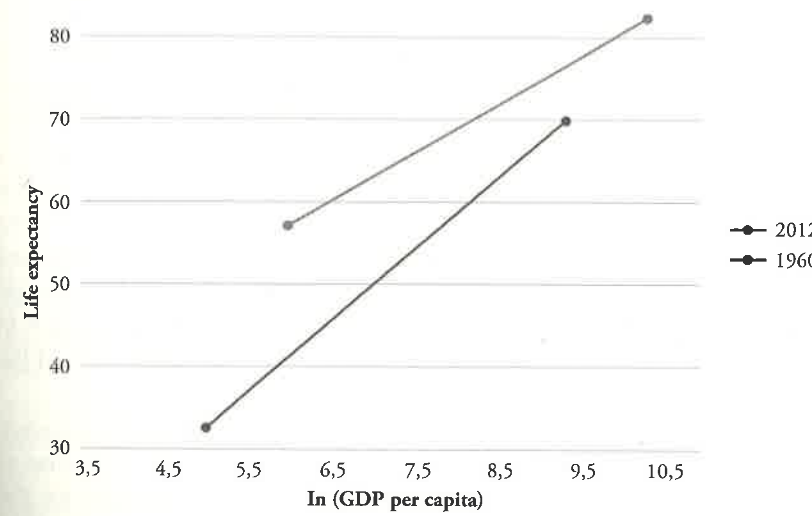

Figure 3.5
Stylized Preston Curves (unweighted estimates)

growth rates of either life expectancy or income in Japan, but my goal is to create a stylized picture. I assume that the Preston Curve for 1960 went through this point, and I give it a slope equal to the 1960 analogue of the value used for 2012. I complete one version of this picture for the unweighted estimates of the Preston Curve and one for the weighted estimates.

The figures clarify that contrary to the simple description of the increase in life expectancy in a country being due to either movement along the Preston Curve or a shift of the Preston Curve, a third force is at work, namely the change in the slope of the Preston Curve. This change is considerably more notable in the case of the unweighted estimates than in that of the weighted estimates. Further, by construction, the effect of changing the slope of the Preston Curve becomes more significant, the poorer a country is, because the curve is anchored using the assumption that rich countries experience income and life expectancy growth according to the description of Oeppen and Vaupel (2002).

These stylized curves can in turn be used to address the question with which Preston started—specifically, how much of the growth of life expectancy is due to time effects (shift or rotation of the Preston Curve) and how much is due to income growth within a country? Specifically, in tables 3.2 and 3.3, I proceed as follows: I consider a set of hypothetical countries that have incomes per capita in 1960 that are 100%, 50%, 20%, 10%, 5%, and 2% of the level assigned to Japan. For each of these countries, I calculate life expectancy according to the estimated 1960 Preston Curve. I then consider two counterfactuals. First, I consider the case in which income per capita in each country grew at a rate of two percent but in which the Preston Curve did not shift from its 1960 position. Second, I consider the case in which income per capita in each country did not change but in which the Preston Curve shifted from its 1960 position to its 2012 position. Finally, I calculate what life expectancy would be if there were both income growth and a shift in the Preston Curve. I conduct the exercise for both the weighted and the unweighted estimates of the Preston Curve.

The bottom part of the table shows the change in life expectancy, relative to the 1960 value, for each of these three scenarios. Because of an interaction effect, the sum of the increases in the columns for the two counterfactuals does not exactly equal the change in life expectancy when both income growth and the curve shift are considered.

These tables demonstrate that Preston's finding of curve shifts being the dominant driver of life expectancy gains holds up relatively well, particularly for very poor countries. For a country with income per capita equal to 2% of the leader, the gain in life expectancy from the shift in the Preston Curve is almost twice as large as the gain from income growth, using either the weighted or the unweighted estimates.

Table 3.2

(A) Predicted Life Expectancy Using Weighted Preston Curve

Income Relative to Japan in 1960	Predicted Life Expectancy in 1960	Life Expectancy if Only Income Growth	Expectancy if Only Shift in Preston Curve	Predicted Life Expectancy in 2012
100%	69.8	76.5	77.0	82.8
50%	65.4	72.0	73.1	78.9
20%	59.5	66.1	68.0	73.8
10%	55.0	61.7	64.1	69.9
5%	50.6	57.3	60.2	66.1
2%	44.7	51.4	55.1	60.9

(B) Decomposition of Increases in Life Expectancy Using Weighted Preston Curve

Income Relative to Japan in 1960	Gain in Life Expectancy if Only Income Growth	Gain in Life Expectancy if Only Shift in Preston Curve	Total Gain in Life Expectancy in 2012
100%	6.7	7.2	13.0
50%	6.7	7.8	13.6
20%	6.7	8.5	14.3
10%	6.7	9.1	14.9
5%	6.7	9.6	15.5
2%	6.7	10.4	16.2

For countries closer to the middle of the income distribution, the relative effect of curve shifts is somewhat smaller. For example, using the unweighted estimates, a country with 20% of the leader's income would have experienced an increase in life expectancy of 8.9 years from income growth alone, versus 11.2 years from the shift in the Preston Curve alone; if we simply divide each of these by the sum, it implies that 44% of the life expectancy gain was due to income growth (doing a similar exercise with the weighted estimates happens to also yield 44%). For the leading country, the exercise reveals that income growth and the shift of the Preston Curve are roughly of equal import.

There are several possible reasons why this exercise produces a different conclusion from the original Preston analysis. First, Preston was examining a period of very rapid health improvement, especially in the developing countries. Specifically, the period that he examined (1930–1960) largely overlapped with the international epidemiological

Table 3.3
(A) Predicted Life Expectancy Using Unweighted Preston Curve

Income Relative to Japan in 1960	Predicted Life Expectancy in 1960	Life Expectancy if Only Income Growth	Life Expectancy if Only Shift in Preston Curve	Predicted Life Expectancy in 2012
100%	69.8	78.7	76.7	82.8
50%	63.8	72.8	72.6	78.7
20%	56.0	64.9	67.2	73.3
10%	50.0	59.0	63.1	69.3
5%	44.1	53.0	59.1	65.2
2%	36.2	45.1	53.7	59.8

(B) Decomposition of Increases in Life Expectancy Using Unweighted Preston Curve

Income Relative to Japan in 1960	Gain in Life Expectancy if Only Income Growth	Gain in Life Expectancy if Only Shift in Preston Curve	Total Gain in Life Expectancy in 2012
100%	8.9	6.9	13.0
50%	8.9	8.8	14.9
20%	8.9	11.2	17.4
10%	8.9	13.1	19.2
5%	8.9	15.0	21.1
2%	8.9	17.5	23.6

transition, during which there was a rapid transfer of health technology from rich to poor countries, motivated by both humanitarian and political reasons (see Acemoglu and Johnson 2007). Second, the calculation presented is by design far more stylized than Preston's. Specifically, I do not look at income and health improvement in actual countries but rather examine hypothetical countries in which income growth is fixed (at 2% per year), and life expectancy is determined exactly by the estimated Preston Curve. Finally, the sample of countries used and the underlying data are not the same.

3.2 Causal Links between Income and Health

Seeing two variables that seem to move together, economists are naturally drawn to ask about causality. Regarding the cross-sectional relation between income and health (the Preston Curve), it is natural to think of such a causal relation in the statistical sense, with there being some underlying structural equations, error terms, and so on. Regarding the data on cutting-edge countries, the statistical interpretation is a bit strained because there are so few observations. However, the broader notion of causality remains relevant.

In the following sections, I discuss possible causality running from income to health and from health to income. I also question whether causality operating in cross section is the same as causality operating over time.

3.2.1 Causality Running from Income to Health

Rising income could have led to improved health through several different channels. One such channel is directly through improvements in the standard of living. In the context of historical improvements in standards of living, this point is most closely associated with the work of McKeown (1976) and Fogel (1993). Fogel, for example, concludes that changes in height and weight alone (due to better nutrition) explain 90% of the reduction in French crude death rates between 1785 and 1870 and a further 50% of the reduction between 1870 and 1975.

The case for the importance of rising standards of living in promoting health improvement is severely diminished for the twentieth century, particularly its second half. The battle against infectious diseases was largely won by the middle of the century, and from that point onward, the availability of antibiotics meant that good nutrition (or certainly, continuing improvements in nutrition) was not nearly as important as it had historically been in reducing deaths. Indeed, it was in this period that too much nutrition became a health problem on a wide scale for the first time in human history. Similarly, the individual contributors to mortality reduction since the middle of the twentieth century (smoking reduction, effective control of high blood pressure) are related to the advancement of knowledge rather than increases in income.

A second channel by which rising income could have led to improved health is through explicit spending directed toward this end, both public health expenditures and spending on medical care. Richer countries (or people) would automatically spend more on health maintenance over time if such spending were a constant fraction of their income. However, there is a second effect, as stressed by Hall and Jones (2007). In standard models of utility, the marginal utility of consumption falls in the level of consumption. By contrast, the utility gain from living (and consuming in) an extra period is assumed to be invariant to the number of periods one lives (except for a small effect from time discounting). This implies that if there is scope to spend money to increase longevity, the ratio of such spending to regular consumption will rise with income. This is Hall and Jones' explanation for the fact that health spending as a fraction of US income rose from 5.2% in 1950 to 15.4% in 2000.

However, the extent to which higher spending was important in contributing to rising life expectancy during the last half century is debatable. Among rich countries, large variations in health spending per capita are largely unrelated to health outcomes. Most significantly, the entire Preston Curve shifting upward over time suggests, just as Preston concluded in the case of developing countries, that there are nonincome-related factors driving mortality improvements throughout the cross-country income distribution. Hall and Jones do their own calculation of the fraction of age-specific mortality reduction due to increased spending in the United States between 1950 and 2000 (from both rising income and a rising share of income devoted to health). They estimate a 32% reduction.

Another reason to doubt the importance of income growth as a contributor to health improvement is direct observation of the sources of improved health. In rich countries, one of the most important sources has been a decline in smoking, something that is well explained by rising knowledge of its health effects rather than increased income. More generally, technological progress in health production, rather than growth of income, appears to be the best explanation for rising health.

Turning to the cross-sectional relationship between income and health, there is again good reason to doubt that causality from income to health is very strong. The short-run relationship between income growth and life expectancy changes is weak, and in particular, one sees in the data episodes of rapid income growth unaccompanied by large changes in health and vice versa. For example, Bolivia, Honduras, and Nicaragua all experienced gains in life expectancy of approximately twenty years during periods of slow or negative income growth (Soares 2007). China's remarkable improvement in infant mortality occurred before its economic growth took off in 1980. Similarly, the acceleration in economic growth in India following economic reforms in the early 1990s was accompanied by a slowdown in the rate of decline in infant mortality (Cutler, Deaton, and Lleras-Muney 2006). Easterly (1999)

examines cross-country data for the period between 1960 and 1990 on income per capita and a number of health indicators. In decadal data, income growth is linked to lower life expectancy, while the relationship between income growth and infant mortality has the expected sign. Income growth is also positively related, though only sometimes statistically significantly, with observable inputs into health, such as calorie intake, physicians per capita, and access to clean drinking water.

3.2.2 Causality from Health to Income

3.2.2.1 Direct Productivity Effects

Measuring the effect of health improvements requires some estimate of the structural effect of health on output. Weil (2007) creates such a measure by examining microeconomic data and taking advantage of a natural experiment, specifically, exogenous variation in uterine nutrition among monozygotic twins. Within pairs of monozygotic twins, there are significant variations in birth weight, reflecting differences in intrauterine nutrition due to the location of the fetuses within the womb. Behrman and Rosenzweig (2004) analyze a sample of female monozygotic twins from the United States. Within this group, the average absolute gap in birth weight is 10.5 ounces, compared to a mean birth weight of 90.2 ounces. Behrman and Rosenzweig regress the within-pair difference in log wages, adult height, and education on the difference in fetal growth (measured in ounces per week of gestation). Their estimate is that a one-unit difference in fetal growth leads to a gap of 0.190 (standard error of 0.077) in log wages, 3.76 (0.43) centimeters in adult height, and 0.657 (0.211) years of schooling. Dividing the estimated effect of fetal growth on log wages by the estimated effect of fetal growth on height gives a TSLS estimate of the return to height of 5.1% per centimeter. This return includes the effect of improved health in raising education. Adjusting to eliminate this channel (see Weil 2007 for details) yields an estimated effect of health as proxied by height on wages, holding education constant, of 3.3% per centimeter. A similar calculation using data on Norwegian twins from Black and Salvanes (2007) yields an estimate of the same effect of 3.5% per centimeter. In the calculations that follow, I use the average of these two estimates, 3.4% per centimeter.[4]

This estimate of the return to height can be applied to historical data. In the typical developed country, the rise in adult height over the past 200 years has been roughly 10 centimeters. Weil's estimate of the return to height thus implies that labor input per worker increased by a factor of 1.385 (in the steady state of a standard growth model, this will also be the effect on output per worker). In annual growth terms (again, in steady state), the contribution from this direct productivity (and associated endogenous accumulation of physical capital) would be only 0.16% per year.

A somewhat more precise version of this calculation can be performed by following Weil (2007) and translating the return to height into terms related to mortality. In that paper, I create a mapping from changes over time in height to changes over time in the Adult Survival Rate (ASR), which is defined as the probability of living from age 15 to age 60 using the current life table. This measure (rather than life expectancy) is chosen because it pertains most clearly to the people who are in the labor force and producing output, with the implicit assumption that the probability of death is proportional to general ill health. Using the data on ASR and height from a panel of countries and going back almost two centuries, I estimate that a change in the ASR of 0.1 is associated with a change in height of 1.92 centimeters. This implies a return to ASR of 0.653, which in turns implies that an increase in ASR by 0.1 would raise labor input per worker by 6.7%.

Using this estimate, in Weil (2007), I then ask how much of the variance in cross-country income can be explained by health. Following Caselli (2005), variance in log output per worker is decomposed into pieces attributable to physical capital, human capital in the form of education, human capital in the form of health, and a productivity residual. The variance in log output per worker is equal to the sum of the variances of these component terms, along with a full set of covariances. I then calculate the reduction in the variance of the log output per worker that would result from eliminating health gaps among countries. The calculation yields a value of 9.9 percent. I also calculate the contribution of health to the 90:10 income ratio (i.e., the ratio of income per capita in the 90th percentile country to that in the 10th percentile country). In the raw data, the ratio is 20.5. Eliminating health gaps, the ratio would fall to 17.9, with the large majority of that reduction coming from a fall in the 50:10 income ratio.

These results reveal that health is a significant contributor to cross-country income differences, but it is not of overwhelming importance. For comparison, the effects of health estimated here are a little more than one-third as large as the contribution of human capital in the form of education to cross-country income variance.

3.2.2.2 Incentives to Accumulate Education

Reduced mortality could affect human capital accumulation by giving people an incentive to stay in school longer, since such investments could be amortized over a longer working life. This idea has a long pedigree in economics (see Kalemli-Ozcan, Ryder, and Weil 2000 and Hazan 2009, the latter of whom traces the mechanism to Ben-Porath 1967).[5]

To assess the potential size of this effect, I consider a simple model in which individual earnings are proportional to human capital, h, which is in turn a function of years of schooling, $s{:}h = f(s)$. I abstract from trend growth in wages and the method by which schooling is financed as well as the risk associated with mortality and simply assume that schooling is chosen to maximize the expected present discounted value of lifetime earnings. Further, for simplicity, I assume that the only cost of schooling is the opportunity cost of foregone wages. The value of s is chosen to maximize

$$\int_s^{\infty} S(a) f(s) e^{-ra}\, da,$$

where $S(a)$ is the probability of survival to age a. For the $f(s)$ function, I use the specification from Klenow and Bils (2000):

$$f(s) = \frac{\Theta}{1-\Psi} s^{1-\Psi}$$

Based on cross-country data on the Mincerian return to schooling, they estimate values of $\Theta = 0.32$ and $\Psi = 0.58$. I take age zero (the first age at which schooling is possible) to be five. To match the example studied by Hazan, I start by using data on survival (the $S(a)$ function) from age five for the cohort of males born in the United States in 1850, when life expectancy at age five was 52.5 years.[6] Hazan reports that this cohort received an average of 8.7 years of schooling. I choose the real interest rate so that optimal schooling matches this value. The implied value of r is 8.7%, which might be considered high. However, given that human capital investment carries risk and that the discount rate applied may reflect credit market imperfections, I do not consider this value unreasonable.

To assess the effect of declining mortality, I hold the other parameters constant and change the $S(a)$ function to match that of the cohort born in 1930, for which life expectancy at age five was 66.7 years. In this case, optimal schooling rises to 9.6 years.[7]

The economic effect of this increased schooling can be evaluated using the standard Mincerian approach. The increase in schooling of 0.9 years over a period of 80 years, or 0.0113 years of education per calendar, is multiplied by an assumed return to schooling of 7% per year, yielding an increase in average human capital per worker of 0.078 percent per year. Assuming again that the country is on a balanced growth path, the increase in output per worker, inclusive of the induced increase in physical capital per worker if the saving rate is fixed, is this same 0.078 percent per year, which is about half the size of the direct productivity effect. Thus, adding this education effect to the direct productivity effect still yields a relatively modest effect of health improvement on income.

3.2.3 Causality in Neither Direction?

An odd conclusion from the previous two sections (discussed at greater length in Weil 2014 and Weil 2015) is that there is not very strong causality in either direction, either from income to health or from health to income. The question then becomes, what does explain the strong statistical relationship between the two (both the cross-sectional Preston Curve and the fact that both began growing during the same historical era)?

The question can be made more precise by constructing a very simple statistical framework. Specifically, we can think of income, y, and health, h, being simultaneously determined:

$$h_i = \alpha y_i + \mu_i \tag{1}$$

$$y_i = \beta h_i + \epsilon_i \tag{2}$$

Discussion of the relationship between health and income can then be conducted in terms of the slope coefficients as well as the variances and co-variances of the two error terms. The previous sections argue that both α and β are relatively close to zero. If this is correct, then the observed correlation between health and income must result from the errors in these two equations being highly correlated. Such an idea should not be surprising. The first error term, μ, contains everything other than income that affects health, while the second, ϵ, contains everything other than health that affects income. In both cases, there are numerous omitted variables that we could imagine being part of these error terms. Further, there are many omitted variables that we would expect to enter both equations. Among these omitted variables, one could list institutions, geography, history, and so on. With so many common elements, we should not be surprised that the error terms are so highly correlated.

However, the statement that the correlation of errors is driven by common omitted variables actually has more economic content than it might seem. Consider a set of factors, X, that affects both error terms, and, for simplicity, assume that these effects are additive and linear. The equations for the two error terms are then

$$\epsilon_i = \sum_j \gamma_j X_{i,j} \text{ and} \tag{3}$$

$$\mu_i = \sum_j \delta_j X_{i,j}. \tag{4}$$

The correlation between ϵ and μ will be determined by the variances and covariances of the X terms, as well as the two sets of parameters, γ and δ. Roughly speaking, there are two possible situations under which ϵ and μ will be very highly correlated. The first possibility is that a single X variable explains most of the variance in both ϵ and μ. In this case, it is easy to see why the two will be correlated. The other possibility is that there are several X variables that contribute to the variances of both ϵ and μ and that ratios of the parameters relevant to these variables in equations (3) and (4) are roughly equal, that is $\frac{\gamma_1}{\delta_1} \approx \frac{\gamma_2}{\delta_2} \approx \ldots$.

The second of these theories seems unlikely. It would be a strange coincidence if two or more of the important omitted factors were affecting health and income with the same ratio. It is possible for one X variable to explain most of the variance in ϵ while a different X variable explains most of the variation in μ, and these two X variables happen to be very highly correlated. However, possibilities such as this are not very plausible either. By contrast, the first theory mentioned previously, that a single omitted variable explains most of the variation in both income and health, does not require any great coincidence.

The question then becomes: what could that omitted variable be? A natural candidate is institutions in their broadest sense. In the case of health, Deaton (2013) argues that differences in health outcomes across countries are highly dependent on governmental capacity to organize public health measures such as clean water and sanitation, to effectively mount public information campaigns to encourage private health behaviors such as hand-washing and condom use, to regulate harmful behavior such as smoking, to provide community-level health services, and to supervise private health providers. Deaton believes these measures, rather than income, explain differences in mortality. In the case of income, the argument that institutions are the dominant determinant of cross-country variation has been forcefully laid out by Acemoglu and Robinson (2012) as well as many other researchers.

To explain the comovement of health and income in the time dimension, at least in the most advanced countries, this sort of institutional explanation is probably not viable. At a practical level, measuring changes in institutions over time is simply more challenging than measuring institutional differences across countries. The general tenor of the development accounting literature (see Caselli 2005) is that productivity is a function of technology level and institutional quality. However, when we do growth accounting, we derive the Solow residual and assume that it is technological progress. Thus, there is a presumption (or at least a widely made assumption) that institutions governing output production are of roughly constant quality over time. At least, this is the assumption we make for rich countries. One would not do the same exercise in China, for example, since it clearly experienced change in institutional quality, rather than technological in the narrow sense, starting around 1980, which explains its subsequent growth. Further, it is clear that institutions governing the production of health do change in cutting-edge countries. One needs only to consider the Affordable Care Act (Obamacare) in the United States, which, depending on who you ask, will have great positive or negative effects on both health and output. Similarly, within the United States, one could cite the rise of environmental regulation, the Pure Food and Drug Act (1906), and the creation of Medicare in 1965 as institutional changes that affected health. As a result, if anything, institutional change appears to be a poor candidate for producing the matched improvements of health and income that we have observed.

A more promising source of these comovements is technological change, with the dimensions of technology applicable to health and output production synchronized. This concurrent change could occur because the growth rates of both of these specific applied technologies are governed by the growth of a more general scientific understanding. A believer in this theory could also point to many ways in which general productive technology has spilled into health-care provision, such as the digitization of medical records, the use of high-speed computers in gene sequencing, and so on. (However, it is harder to think of cases in which medical technology has spilled over to production in nonhealth-related areas).

3.3 Conclusion

In this essay, I explore the empirical relationship between income per capita and life expectancy, both across countries and over time. I focus in particular on the slope of the average relationship between these two measures, expressed as years of life expectancy per log unit of GDP per capita. In the cross section, where this relationship is called the Preston Curve, data for 2012 traces out a slope of roughly 5.7 years of life expectancy per log unit of GDP per capita. The slope of this relationship is flatter for 2012 than for 1960, although exactly how much flatter depends on how one assigns weights to countries with different populations. Over time, among the richest and healthiest countries, the relationship between income growth and health improvement is roughly 12.5 years of life expectancy per log unit of GDP per capita. The fact that this slope is higher than the cross-sectional slope is testament to a phenomenon first noted by Preston (1975)—the Preston Curve is shifting upward over time. This upward shift is attributable to accumulation of health knowledge.

Using estimates of the slope and position of the Preston Curve in 1960 and 2012, I conduct a decomposition of increases in life expectancy that occurred over this period into pieces due to increased income (movement along the Preston Curve) and shifts in the Preston Curve over time. The approach I take is somewhat more stylized than the one originally used by Preston, and the period and country sample also differ. Because the slope of the Preston Curve changes over time, the relative importance of curve shifts and income growth depends on a country's place in the income distribution. For a country with 20% of the income per capita of the leaders, my estimate is that 44% of life expectancy gains were attributable to income growth and 56% to the shifting of the Preston Curve. For countries at the cutting edge, life expectancy growth was 50% due to each of these factors. By contrast, Preston's estimate was that, in developing countries, no more than one-third of life expectancy gain was due to income growth.

The final part of the essay addresses the question of the extent to which the cross-sectional relationship between income and life expectancy is shaped by causality running from income to health, and how much it is shaped by causality running from health to income. I argue that causality in both directions is, in fact, rather weak. Rising GDP per capita in a country, ceteris paribus, does relatively little to raise life expectancy. Similarly, although better health does lead to higher income, the effect is not large. These observations lead me to conclude that most of the observed cross-sectional correlation is induced by the operation of a third factor, specifically, cross-country variation in the quality of institutions. Good institutions raise income directly, holding health constant, and similarly improve health directly, holding income constant.

Notes

1. Data from Cuter and Meara (2004). As they point out, this phenomenon does not result from the fact that rates of mortality decline (relative to the level of mortality) at young ages have slowed down but rather from the fact that mortality rates at young ages are now so low that a percentage decline in rates leads to a very small increase in the probability of survival.

2. Centers for Disease Control and Prevention, "Leading Causes of Death," 1900 to 1998, http://www.cdc.gov/nchs/data/dvs/lead1900_98.pdf.

3. It is worth mentioning that in the case of income, the growth rate is constant, while in the case of life expectancy, the time derivative is constant. I note that this is how the empirical regularities are found in the data, without having a particular theory for why it should be so.

4. Weil (2007) applies these estimates of the return to height and ASR to cross-country data. However, the underlying microeconomic evidence comes from rich countries, where the mix of labor market activities and economic structure are presumably not well representative of those in the sample as a whole. Thus, this application to growth among cutting-edge countries is actually more appropriate than the original exercise.

5. Hazan criticizes the Ben-Porath mechanism, arguing that for an increase in survival to induce a rise in schooling it must also induce a rise in lifetime labor supply. In his paper, Hazan measures expected total working hours (ETWH) over the lifetime for cohorts of American men born between 1850 and 1970. In addition to mortality, ETWH is affected by labor supply along both the extensive margin (working or not working) and the intensive margin (hours per week). He shows that declines in weekly hours, along with earlier retirement, have more than offset the decline in mortality. For example, ETWH at age 20 fell from 112,199 for men born in 1850 to 81,411 for men born in 1930.

Hazan's observation that ETWH have fallen over time is indeed well taken, but it is worth noting that the paper does not actually show that changing mortality did not affect schooling. Rather, it shows that even though falling mortality worked to increase ETWH, other factors more than undid this effect. Presumably, if mortality had not fallen, ETWH would have fallen more than what we observe, and schooling would have risen less. Another critique of Hazan is that, as pointed out by Cervellati and Sunde (2013), reduced labor supply on the intensive margin (i.e., fewer hours worked per year) decreases the opportunity cost of schooling as well as the benefits to additional years of education.

6. I am grateful to Moshe Hazan for sharing this data.

7. In fact, average years of schooling for this cohort was 13.3. Thus, the pure mortality effect explains roughly one-fifth of the actual increase in schooling that took place during this period. This exercise shows that reduced mortality over the range found in historical data should have some effect on schooling, but we would not expect it to be the dominant explanation.

References

Acemoglu, D., and S. Johnson (2007): "Disease and Development: The Effect of Life Expectancy on Economic Growth", *Journal of Political Economy*, 115, 925–985.

Acemoglu, D., and J. Robinson (2012): *Why Nations Fail: The Origins of Power, Prosperity, and Poverty*, New York: Crown Business.

Behrman, J. R., and M. R. Rosenzweig (2004): "Returns to Birthweight", *Review of Economics and Statistics*, 86, 586–601.

Ben-Porath, Y. (1967): "The Production of Human Capital and the Life Cycle of Earnings", *Journal of Political Economy*, 75, 352–365.

Black, Sandra E., P. J. Devereux, and K. G. Salvanes (2007): "From the Cradle to the Labor Market? The Effect of Birth Weight on Adult Outcomes", *Quarterly Journal of Economics*, 122, 409–439.

Caselli, F. (2005): "Accounting for Cross-Country Income Differences", in *Handbook of Economic Growth*, eds. P. Aghion and S. N. Durlauf, North-Holland, volume 1, 679–742.

Cervellati, M and Sunde, U (2013) "Life Expectancy, Schooling, and Lifetime Labor Supply: Theory and Evidence Revisited" *Econometrica* 81(5), 2055–2086.

Cuter, D., and E. Meara (2004): "Changes in the Age Distribution of Mortality over the Twentieth Century", in *Perspectives on the Economics of Aging*, ed. D. A. Wise, University of Chicago Press, 333–365.

Cutler, D., A. Deaton, and A. Lleras-Muney (2006): "The Determinants of Mortality", *Journal of Economic Perspectives*, 20, 97–120.

de la Croix, D., and O. Licandro (2015): "The Longevity of Famous People from Hammurabi to Einstein", *Journal of Economic Growth*, 20, 263–303.

Deaton, A. (2013): *The Great Escape: Health, Wealth, and the Origins of Inequality*, Princeton University Press.

Easterly, W. (1999): "Life During Growth", *Journal of Economic Growth*, 4, 239–275.

Fogel, R. W. (1993): "New Findings on Secular Trends in Nutrition and Mortality", in *Handbook of Population and Family Economics*, eds. M. R. Rosenzweig and O. Stark, Elsevier Science Publishers, volume 1A, 433–481.

Galor, O., and D. N. Weil (2000): "Population, Technology, and Growth: From Malthusian Stagnation to the Demographic Transition and Beyond", *American Economic Review*, 90, 806–828.

Hall, R. E., and C. I. Jones (2007): "The Value of Life and the Rise in Health Spending", *Quarterly Journal of Economics*, 112, 39–72.

Hazan, M. (2009): "Longevity and Lifetime Labor Supply: Evidence and Implications", *Econometrica*, 6, 1829–1863.

Jones, C. (2002): "Sources of U.S. Economic Growth in a World of Ideas", *American Economic Review*, 92, 220–239.

Kalemli-Ozcan, S., H. E. Ryder, and D. N. Weil (2000): "Mortality Decline, Human Capital Investment, and Economic Growth", *Journal of Development Economics*, 62, 1–23.

KEYNES, J. M. (1930) "Economic Possibilities for our Grandchildren", *The Nation a* *Athenaeum*, 1930

KEYNES, J. M. (1972): "Economic Possibilities for our Grandchildren", in *The Collect* *Writings of John Maynard Keynes, volume 9, Essays in Persuasion*, London: Macmillan.

KLENOW, P. J., AND M. BILS (2000): "Does Schooling Cause Growth?" *American Econom* *Review*, 90, 1160–1183.

LUCAS, R. E. (2000): "Some Macroeconomics for the 21st Century", *Journal of Econom* *Perspectives*, 14, 159–168.

MADDISON, A. (2005): *Growth and Interaction in the World Economy: The Roots of Modernit* Washington, DC: The AEI Press.

MALTHUS, T. (1798): "An Essay on the Principle of Population", London: J. Johnson.

MCKEOWN, T. (1976): *The Modern Rise of Populations*, New York: Academic Press.

OEPPEN, J., AND J. W. VAUPEL (2002): "Broken Limits to Life Expectancy", *Science*, 296, 1029–1031.

PELTZMAN, S. (2009): "Mortality Inequality", *Journal of Economic Perspectives*, 23, 175–190.

PRESTON, S. H. (1975): "The Changing Relation between Mortality and Level of Eco- nomic Development", *Population Studies*, 29, 231–248.

SOARES, R. R. (2007): "On the Determinants of Mortality Reductions in the Developing World", *Population and Development Review*, 33, 248–287.

WEIL, D. N. (2007): "Accounting for the Effect of Health on Economic Growth", *Quarterly Journal of Economics*, 122, 1265–1306.

——— (2014): "Health and Economic Growth", in *Handbook of Economic Growth*, eds. P. Aghion and S. N. Durlauf, North Holland, volume 2, 623–682.

——— (2015): "A Review of Angus Deaton's The Great Escape: Health, Wealth, and the Origins of Inequality", *Journal of Economic Literature*, 53, 102–114.

4 Did Longer Lives Buy Economic Growth? From Malthus to Lucas and Ben-Porath

David de la Croix*

4.1 Introduction

To what extent long lives matter for growth is a topic that has been investigated both theoretically and empirically in history and in contemporary data. Accordingly, I first provide empirical evidence indicating that improvements in life expectancy occurred before the take-off to modern growth. Establishing the precedence of longevity over growth is one argument in favor of causality. After a short section on measurement, I discuss two mechanisms, the contact time effect and the incentive effect, through which longevity may impact growth, both in the past and today, and their quantitative significance.

4.2 Early Longevity Increases

At the time of this publication, I would claim that the precedence of longevity improvements for the elite during the Industrial Revolution is firmly established. That longevity increased in the seventeenth and eighteenth centuries was known by historian demographers on the basis of local evidence and specific social groups. For example, Hollingsworth (1977) builds mortality tables for British peers sampled from genealogical data. Vandenbroucke (1985) provides vital statistics for the knights of the Order of the Golden Fleece, which was founded in 1430 by the Dukes of Burgundy and continued by the Habsburg rulers, the kings of Spain, and the Austrian emperors. In both samples, mortality reduction for nobility occurred as early as during the

*IRES, Université Catholique de Louvain, Place Montesquieu 3, B-1348 Louvain-la-Neuve, Belgium. E-mail: david.delacroix@uclouvain.be