PROGRAM ON THE GLOBAL DEMOGRAPHY OF AGING

Working Paper Series

Disease and Development Revisited

David E. Bloom, David Canning, Günther Fink

June 2009

Disease and Development Revisited[1]

David E. Bloom
David Canning
Günther Fink

Harvard School of Public Health

June 2009

Abstract

In a recent paper, Acemoglu and Johnson (2007) argue that the large increases in population health witnessed in the 20[th] century may have lowered income levels. We argue that this result depends crucially on their assumption that initial health and income do not affect subsequent economic growth. Using their data we reject this assumption in favor of a model of conditional convergence, with income adjusting to its steady state over time. We show that, allowing for conditional convergence, exogenous improvements in health due to technical advances associated with the epidemiological transition appear to have increased income levels.

1. Introduction

In a recent paper, Acemoglu and Johnson (2007) present results indicating that increases in population health, as measured by higher life expectancy, are negatively correlated with economic growth in cross-country panel data. This appears to contradict much of the preceding literature, surveyed in Weil (2007) and Bloom, Canning, and Sevilla (2004), which generally finds that countries with better health achieve higher rates of economic growth.

Acemoglu and Johnson express three concerns about the findings of previous studies linking population health to the growth of income per capita. The first concern involves the possibility that the estimated positive coefficient of health in an empirical growth equation is biased upward by the omission of some variable that is positively associated with both good health and high income growth. As the authors argue, countries with low life expectancy may be disadvantaged in other ways, so that the negative effects of other country-specific factors might be falsely attributed to ill health. The second concern focuses on reverse causality from economic growth to population health, since health depends on health inputs that are more affordable in richer countries. The third issue is a general equilibrium effect of health on economic growth via lower mortality and larger population numbers, an effect that will not be captured by individual-level studies.

Given these concerns, Acemoglu and Johnson argue that the appropriate way to identify the total effect of health on economic growth is empirical work at the aggregate level with valid instrumentation to address the critical endogeneity issue. The instrument they propose depends on exogenous shocks to national health generated by improvements in health technology. Technological innovations in prevention and treatment occur for different diseases at different times.[2] They also affect health differently in each country since the prevalence of each disease varies across countries.

---

[2] We follow Acemoglu and Johnson in regarding disease specific technological advances as exogenous. However a case could be made that research has focused on the diseases of rich countries, rather than poor countries, making technological advances endogenous.

The issues raised by Acemoglu and Johnson are indisputably central to understanding the health-income links. However, we argue that the discrepancy between the Acemoglu and Johnson results and those reported in earlier literature is not a consequence of addressing these three issues via a new instrument, but rather of the way in which the dynamic relationship between health and income is modeled and estimated. Acemoglu and Johnson specify a formulation in which current income is a function of, and adjusts instantaneously to, current health.

There are both theoretical and empirical reasons to expect income to adjust non-instantaneously to changes in population health. The standard partial adjustment model that lies behind the conditional convergence framework for modeling economic growth assumes that income converges slowly its steady-state level; a convergence rate of about 2% per year is a common estimate (Barro and Sala-i-Martin 2004; Durlauf, Johnson et al. 2005). It is likely that, when health improves, income adjusts slowly to the new steady state. There is growing evidence that physical and cognitive development are closely linked to health and nutrition *in utero* and during the first few years of life. Health human capital may consequently be affected by conditions in the early years of childhood, while the returns to this capital are realized mainly during a person's adult life (Heckman 2007). For example, Barker (1992) and Crimmins and Finch (2006) find that early childhood infections have large effects on adult health. Behrman and Rosenzweig (2004) find that development *in utero*, as measured by birth weight, has long-term economic effects on adult productivity and earnings. Bleakley (2003; 2006) and Kremer and Miguel (2004) find that early childhood exposure to disease has effects on educational outcomes while (Hoddinott, Maluccio et al. 2008) find that early childhood nutrition has substantial effects on adult income. Case, Fertig, and Paxson (2005) and Schultz (2005) provide further evidence of a link between health in childhood and later economic success.

This empirical evidence suggests that health has a positive effect on income in the long term. We are then left with the question why Acemoglu and Johnson find such a strong negative effect of health improvements on growth in per capita income. As we show in the following section, the negative correlation between changes in income and changes in life expectancy found by Acemoglu and Johnson has little to do with instrumentation or the relatively small sample of 47 countries chosen by the authors. The reason this negative correlation emerges is simply that

countries with good health in 1940 were the ones that benefited least from subsequent global health technology improvements, but are also the countries whose economies have grown fastest since 1940. Acemoglu and Johnson ignore the effect of initial conditions on economic growth and limit their attention to the unconditional negative correlation between health gains and economic growth. Since there is strong convergence in health, countries with relatively good initial health conditions have relatively smaller subsequent health gains (Cutler, Deaton et al. 2006). The unconditional regression of changes in income on changes in life expectancy omits initial conditions as one of the fundamental drivers of subsequent economic growth. Since the correlation between initial health conditions and subsequent changes in health is strongly negative, the resulting estimates display a negative bias. The magnitude of the omitted variable bias is large; once we control for initial conditions, the estimated effect of changes in life expectancy on economic growth changes sign and becomes positive, with large and highly significant coefficients both on initial levels of, and changes in life expectancy.

If the instrumental variable approach chosen by Acemoglu and Johnson were valid, this omitted variable (or selection) bias would be irrelevant. A good instrument essentially mimics a random assignment of health improvements, and thereby allows researchers to identify the causal effect of health on income, independent of other omitted variables. Improvements in health technology during the last century could be regarded as random shocks, and we could regard these innovations as plausible instruments. Major advances in our understanding of infectious diseases and the emergence of cheap and effective prevention and treatment interventions have indeed led to rapid health improvements. Since the initial burden of disease varies across countries, these health improvements provide differential positive shocks to all countries. Acemoglu and Johnson discuss innovations developed for 15 prominent diseases in the early 20[th] century, ranging from measles to tuberculosis to malaria[3]. They take "predicted mortality" to be the actual mortality rate from these diseases in 1940, but once the cure for a specific disease is accessible to a country, the predicted mortality from that disease is taken to be zero. The assumption that the burden of malaria drops to zero after the intervention of DDT might appear odd from a historical

---

[3] The 15 diseases are tuberculosis, malaria, pneumonia, influenza, cholera, typhoid, smallpox, whooping cough, measles, diphtheria, scarlet fever, yellow fever, plague, typhus fever, and dysentery/diarrhea-related diseases.

4

perspective. It makes sense from an econometric perspective, however, as it avoids confounding the exogenous health shock with endogenous domestic demand and implementation factors.

The problem with the instrument constructed by Acemoglu and Johnson is that it assumes the predicted mortality to be exogenous and not affected by contemporaneous income shocks. This implies that the initial mortality rate in 1940 due to each disease must be unaffected by income levels in 1940. We think this is implausible, particularly given that the problem being addressed is endogeneity of life expectancy, which is closely linked to mortality. We discuss this issue in detail in the following for the case of malaria, arguing that malaria mortality in 1940 was endogenous, with richer countries in 1940 being better able to afford vector control of mosquitoes, such as swamp clearance.

To put it in the context of the evaluation literature, the "natural experiment" constructed by Acemoglu and Johnson is flawed because the "treatment group" that received large health gains from technological innovations is fundamentally different from the "control group" that received low health gains. The treatment group, which benefits most from technological advances, had poor health initially (low life expectancy) while the control group had good health (high life expectancy). If initial health matters for subsequent economic growth, this difference between the treatment and control groups will bias their result.

In our empirical analysis we focus on economic growth over the period 1940-2000. We expand on the Acemoglu and Johnson model by allowing initial health and income to affect subsequent economic growth. This can be regarded as adding what may be omitted variables to the model. We also show formally that adding these explanatory variables makes the identifying assumption in the model that disease specific mortality rates 1940 are uncorrelated with the income shock in 2000. This seems to us more plausible than the Acemoglu and Johnson assumption than that 1940 mortality is uncorrelated with the income shock in 1940[4]. By controlling for initial health and income we remove the correlation between the instrument and these previously omitted variables. In this alternative framework, we show that improvements in life expectancy over the

---

[4] We test, and cannot reject, that the income shocks in 1940 and 2000 are uncorrelated.

period 1940-2000, instrumented with the change in predicted mortality due to technological innovations, have a positive effect on the growth of income per capita.

To provide a clearer understanding of the results generated by Acemoglu and Johnson, we first analyze the relationship between health, economic growth and predicted mortality in the data set constructed by Acemoglu and Johnson in Section 2 of this paper. We show that there is a strong positive correlation between initial health and subsequent economic growth. Since there is a pronounced negative correlation between initial health and subsequent improvements in health, a negative correlation between improvements in health and economic growth emerges. These correlations are consistent with the view that good initial health leads to rapid subsequent economic growth. They are, however, also consistent with the view that the correlation of changes in income and health is negative.

In Section 3, we show formally that in the model used by Acemoglu and Johnson the identifying assumption is that current income shocks do not affect current disease-specific mortality rates. We also construct an alternative expanded model, in which the predicted mortality instrument used by Acemoglu and Johnson is valid provided only that future income shocks do not affect current disease-specific mortality. In Section 4 we report parameter estimates based on the two models and show that while the Acemoglu and Johnson model finds a negative effect of health improvements on economic growth, our model finds relatively large and highly significant positive effects of health on income.

2. Patterns of Health and Economic Growth

We begin by examining the 47 data points used in Acemoglu and Johnson's main regressions. The data were provided by the authors and are as described in the original paper (Acemoglu and Johnson 2007). Figure 1 shows the relationship between life expectancy in 1940 and growth in income per capita over the period 1940-2000. The figure reveals a positive relationship: countries with high life expectancy in 1940 had more rapid per capita income growth over the period 1940-2000. Figure 2 shows the relationship between life expectancy in 1940 and the growth rate of life expectancy during 1940-2000. Figure 2 shows that initial life expectancy

almost perfectly predicts subsequent percentage gains in life expectancy, with countries with low initial life expectancy having the largest subsequent gains in life expectancy. This relationship reflects a diffusion of health technologies that have allowed rapid health improvements even in poor countries and overall convergence in life expectancy (Cutler, Deaton et al. 2006), though there is evidence that not all countries are converging to the same steady state (Bloom and Canning, 2008).

Figure 3 shows a negative relationship between growth in life expectancy over the period 1940-2000, and the growth rate of income per capita over the same period[5]. Figure 3 is a visual representation of the Acemoglu and Johnson ordinary least squares estimates of the effect of improvements in life expectancy on economic growth; the countries with above average gains in life expectancy also have below average growth in income per capita over the period. These patterns in the data are not due to the selection of particular countries or a particular time period. In Figures 4-6 we plot the same relationships over the period 1970-2000 when data for 112 countries are available[6]. Data for income per capita are from the Penn World Tables 6.2 (Heston, Summers et al. 2006), while data for life expectancy are from the World Development Indicators 2007 (World Bank 2007). This dataset is similar to that used in most recent studies of the effect of health on economic growth. In Figure 4 we see once again a positive relationship between initial life expectancy (in 1970), and subsequent economic growth (over the period 1970-2000). In Figure 5 we see once again a negative relationship between initial life expectancy (in 1970) and the subsequent change in life expectancy (over the period 1970-2000). Last, Figure 6 shows a negative relationship during the period 1970-2000 between growth in life expectancy and growth in income per capita.

Most of the literature on the relationship between health and growth focuses on the positive link between initial health and subsequent economic growth as seen in Figure 1 and Figure 4. The major innovation of Acemoglu and Johnson is not to instrument health in this relationship, but to

---

[5] Following Acemoglu and Johnson, we use log differences as our measure of growth in life expectancy and income per capita. In Figures 1-6, we convert the long-term growth figures into annual growth rates by dividing the log difference by the number of years between the two observations.
[6] We exclude countries with HIV prevalence rates over 3% from these figures. HIV dominates the improvements made with respect to all other diseases in high-prevalence countries and makes the comparison of aggregate time trends in life expectancy rather difficult.

shift the emphasis to the negative relationship between growth in life expectancy and growth in income per capita as seen in Figure 3 and Figure 6.

However, if the relationships in Figures 1 and 2 are causal, the relationship in Figure 3 becomes difficult to interpret. Suppose high initial life expectancy does lead to economic growth as suggested by Figure 1. Suppose in addition that low initial life expectancy leads to a high rate of growth in life expectancy as suggested by Figure 2. Combining these two mechanisms we would expect to see that countries with high initial life expectancy in 1940 have high rates of economic growth over the period 1940-2000 and relatively low rates of growth of life expectancy over the same period. Failing to take account of both mechanisms can yield a spurious negative correlation between growth in income and growth in life expectancy, even if the true relationship is positive.

Acemoglu and Johnson continue to find a negative coefficient on growth in life expectancy in regressions explaining growth in income per capita even when they instrument the change in life expectancy with predicted mortality change. The predicted mortality variable they construct is the total number of deaths per 100 people living in a given country attributable to 15 diseases that underwent significant advances in prevention or treatment over the period. Predicted mortality is the actual number of deaths prior to the health innovation, and is set to zero thereafter. The countries with the largest declines in predicted mortality are therefore the countries with high initial mortality, which means low initial life expectancy.

Figure 7 plots changes in predicted mortality, the Acemoglu and Johnson instrument, against the log of initial life expectancy in 1940. This graph shows that the countries that experienced the largest health improvements (largest reductions in predicted mortality) are those with the lowest initial life expectancy. This instrument is clearly not a random assignment. Its validity depends on being able to exclude initial income as an explanatory variable for subsequent economic growth. If this exclusion is not possible, the instrument will be correlated with growth via its correlation with initial life expectancy as well as through its correlation with growth in life expectancy.

An intuitive way of illustrating our point about the implausibility of the predicted mortality change instrument is to look at the analysis presented from a natural experiment perspective. Let us define the set of 47 countries analyzed by Acemoglu and Johnson as our sample and hypothetically divide the sample into a treatment group and a control group. The treatment group gets a large shock (i.e., a relatively big increase in life expectancy), while the control group gets only a small shock (i.e., a relatively small increase or even a decline in life expectancy). Following Angrist and Pischke (2009), let us define $D_i = \{0,1\}$ to be the treatment indicator for country i, Δy_i to be growth in income as the outcome variable, and Δy_{i0} to be income growth in the absence of the treatment. The observed difference in growth is then given by

$$\underbrace{E[\Delta y_i \mid D_i = 1] - E[\Delta y_i \mid D_i = 0]}_{\text{Observed difference}} = \underbrace{E[\Delta y_i \mid D_i = 1] - E[\Delta y_{i0} \mid D_i = 1]}_{\text{Treatment effect on the treated}} + \underbrace{E[\Delta y_{i0} \mid D_i = 1] - E[\Delta y_{i0} \mid D_i = 0]}_{\text{Selection bias}}$$

which nicely summarizes the issue at hand. The observed difference between the treatment group (countries with generally low life expectancy in 1940) and the control group (countries with generally high life expectancy in 1940) is the sum of the true treatment effect of improved health and the difference in growth patterns in the absence of the treatment. The observed differences reflect the true treatment effect if there are no differences in expected growth between the treatment and control group in the absence of the health shock. In the sample used in the Acemoglu and Johnson case, this assumption seems quite implausible; treatment is associated with low initial life expectancy which is correlated with low subsequent economic growth, as shown in Figure 1. This implies a negative selection bias. Countries in the control group share initial conditions associated with higher subsequent income growth. This means that countries in the treatment group would have done worse in terms of economic performance in the absence of health improvements than countries in the control group. Observing a negative growth difference between the treatment and control groups does thus not permit any inferences about a possible causal effect of health on income per capita.

Table 1 shows a correlation matrix of the variables plotted in Figures 1-4. Note that the change in predicted mortality is just the negative of predicted mortality in 1940, since predicted mortality in 2000 is uniformly set to zero. The correlation between initial life expectancy and

9

subsequent economic growth is 0.50, while the correlation between predicted mortality in 1940 and life expectancy in 1940 is -0.70. Both correlations are significant at the 1% confidence level. Since initial life expectancy is not accounted for in the empirical specifications implemented by Acemoglu and Johnson, the instrument is correlated with an excluded variable, which will induce inconsistent estimates if this excluded variable is relevant for the model.

3. Model Specification

We now address the issue of the dynamics of the relationship between health and income more formally. The main estimates presented in Acemoglu and Johnson are built around the following empirical specification:

$$y_{it} = \alpha + \pi x_{it} + \beta z_{it} + u_i + \delta_t + \varepsilon_{it} \tag{1}$$

where y_{it} is the outcome variable in country i at time t, x_{it} is log life expectancy, z_{it} represents other explanatory variables, u_i are country fixed effects, δ_t are time dummies and ε_{it} is an error term. Acemoglu and Johnson analyze log total population, log total GDP, log GDP per worker, and log GDP per capita as outcome variables, although the main focus of their study is on the effect of life expectancy on GDP per capita. We focus solely on log GDP per capita as the outcome. In addition, while equation (1) can be run as a panel we focus on the case where we analyze only two time periods, the "long difference", with data from 1940 and 2000. There is an important assumption in equation (1) that health can be reduced to a single measure and can be measured by life expectancy, which may be disputed (Weil 2007; Bleakley 2008).

One critique of Acemoglu and Johnson is that they do not control for a sufficient number of alternative possible explanatory variables z_{it}. They do show that their results are robust to including average institutional quality, and initial population, interacted with a post-intervention dummy, but other variables could be added. In the empirical section below we investigate adding initial life expectancy and initial income, interacted with a post-intervention dummy, as additional explanatory variables.

A second, though as we shall see related, critique concerns the validity of the instrument used by Acemoglu and Johnson. Acemoglu and Johnson argue that life expectancy in equation **(1)** is not exogenous, since income is likely to affect health, so that estimating this relationship by ordinary least squares (OLS) will yield biased results. They propose an instrument, predicted mortality, to avoid this endogeneity problem. Predicted mortality M_{it}^I (mortality from diseases with a health innovation I in the sample period) is given by

$$M_{it}^I = \sum_{d \in D} \left[(1 - I_{dt}) M_{di0} + I_{dt} M_{dFt} \right] \tag{2}$$

where M_{di0} is the mortality rate from disease d in country i at time 0, which is taken to be 1940, and I_{dt} is a dummy taking the value zero prior to technological innovation in disease d and a value of one after the innovation. M_{dFt} denotes the frontier mortality in the disease post intervention, taken to be zero for their baseline instrument. In the "long difference" comparison, where period 0 is 1940 and period 1 is 2000, this reduces to

$$M_{i0}^I = \sum_{d \in D} M_{di0}, \quad M_{i1}^I = 0 \tag{3}$$

Predicted mortality in 1940 is just the sum of observed mortality rates in the diseases under consideration while predicted mortality in 2000 is zero (or at the frontier if the alternative instrument used) since every disease considered has an innovation before 2000. The assumption required for this instrument to be valid is that it is uncorrelated with ε_{it}, the error term in equation (1); that is, $Cov(M_{it}^I, \varepsilon_{it}) = 0$. Given that the instrument is zero for all countries in 2000 this reduces to the assumption that in 1940 the sum of observed mortality rates in the diseases under consideration is uncorrelated with the error term in 1940, that is, $Cov(\sum_{d \in D} M_{di0}, \varepsilon_{i0}) = 0$.

We derive this formally in Appendix I below. While Acemoglu and Johnson argue that this is plausible, our view is that it is not.

The problem being addressed is potential reverse causality from income to life expectancy. Such income effects on health seem likely since higher income can improve individual nutrition and access to health care as well as increase the availability of government funds for public health

measures, all of which are likely to reduce mortality. For the predicted mortality instrument to be valid, Acemoglu and Johnson have to argue that shocks to income in 1940, while affecting mortality and life expectancy in 1940 overall, do not affect mortality rates in the 15 diseases they consider, which seems highly unlikely. We can illustrate this with the example of malaria. Acemoglu and Johnson date technological innovation in malaria to 1947, with the recognition by the Expert Committee on Malaria (1947) that DDT was an effective insecticide against mosquitoes. However, urbanization, vector control based on management and drainage of water sources, the use of insecticides, and access to quinine, had already substantially reduced the burden of malaria in the United States and Italy by 1947, though at a high financial cost, while malaria control efforts in India were hampered by lack of funds (Snowden 2006; Packard 2007). Assuming that the initial burden is exogenous implies assuming that none of the factors explaining the relatively low death burden in Italy and the United States as compared to India in 1940 has any explanatory power for income in 1940, which seems hard to defend.

We can get a clearer picture of the likely bias introduced by this endogeneity by considering equation (1) in first differences. With two time periods only and country fixed effects, it is easy to see that equation (1) can be rewritten as a first difference model given by:

$$\Delta y_i = \delta + \pi \Delta x_i + \beta \Delta z_i + \Delta \varepsilon_i \tag{4}$$

where $\Delta y_i = y_{i1} - y_{i0}$ is the growth rate (log difference) of income per capita between period zero and period one (1940 and 2000), $\Delta x_i = x_{i1} - x_{i0}$ is the growth rate (log difference) of life expectancy, and $\Delta z_i = z_{i1} - z_{i0}$ is the change in other explanatory variables, in country i over the same period. All country-specific time-invariant factors are differenced out, and the constant in the growth equation $\delta = \delta_1 - \delta_0$ represents the change in the time dummies. The error term in the growth equation is $\Delta \varepsilon_i = \varepsilon_{i1} - \varepsilon_{i0}$. We can write the first difference of the predicted mortality instrument as

$$\Delta M_i^I = M_{i1}^I - M_{i0}^I = 0 - M_{i0}^I = -\sum_{d \in D} M_{di0} \tag{5}$$

The change in predicted mortality is just the negative of the predicted mortality in 1940 since mortality in the diseases with new technologies is assumed to go to zero by 2000. The regression (4) using the first differenced predicted mortality (5) gives numerically identical results to running the levels regression (1) for 1940 and 2000 with the level of predicted mortality in levels from equation (3) as the instrument. The problem of endogeneity of the instrument remains. In first differences the condition for instrument validity is

$$Cov(\Delta M_i^I, \Delta \varepsilon_i) = Cov(-\sum_{d \in D} M_{di0,} \varepsilon_{i1} - \varepsilon_{i0}) = 0 \qquad (6)$$

However, if as we have argued above, mortality in the 15 diseases in 1940 is correlated with the income shock in 1940 this covariance will not be zero, unless $\varepsilon_{i1} - \varepsilon_{i0}$ is uncorrelated with ε_{i0}, which is unlikely.[7]

Acemoglu and Johnson argue that variations in M_{it}^I are unrelated to any actions or economic events in the country, which makes predicted mortality a plausible exogenous instrument. We agree with this assessment for the dummy I_{dt}: global technological progress in these diseases is plausibly uncorrelated with each country's economic actions. However, the other component of the instrument is disease-specific mortality rates. Our view is that this component reflects current health conditions in the country, and life expectancy itself, is endogenous to current income. Interacting this endogenous component with the dummy for a technological innovation, or first differencing, does not remove the endogeneity.

We could stop here and conclude that the predicted mortality instrument is unlikely to be valid, either in levels or in first differences. However, we can derive an equation in which predicted mortality is a plausibly valid instrument. We can rewrite equation (4) as

$$\Delta y_i = \delta + \pi \Delta x_i + \beta \Delta z_{it} + (\alpha + u_i + \pi x_{i0} + \beta z_{i0} - y_{i0}) + \varepsilon_{i1} \qquad (7)$$

where we have substituted for the base line income shock ε_{i0} in the differenced error using equation (1). This formulation allows for an effect of change on life expectancy on change in income as in Acemoglu and Johnson. However, we also have a conditional convergence term.

---

[7] $\varepsilon_{i1} - \varepsilon_{i0}$ and ε_{i0} will be uncorrelated only if ε_{it} is a random walk; we test and decisively reject this condition on the data below.

Countries with high life expectancy and low income in 1940, which are therefore below their steady-state income levels, will be expected to have faster economic growth. The advantage of equation (7) is that the condition for the change in predicted mortality to be a valid instrument is now:

$$Cov(\Delta M_i^I, \varepsilon_{i1}) = Cov(-\sum_{d \in D} M_{di0}, \varepsilon_{i1}) = 0 \qquad \textbf{(8)}$$

This seems reasonable. All we require is that mortality rates in the 15 diseases in 1940 are uncorrelated with the shock to income in 2000. While income shocks are likely to be correlated over time it seems plausible that any correlation will have disappeared over a 60-year period. The advantage of equation (7) is that in this formulation the predicted mortality change is a valid instrument. This equation makes explicit the point that the equation estimated by Acemoglu and Johnson, equivalent to equation (7) but without the term in brackets, suffers from omitted variable bias, and it is the correlation with these omitted variables that is likely to make their instrument invalid.

Given the 60-year gap between 1940 and 2000, the assumption that the error terms in levels are uncorrelated is plausible; any disequilibrium in 1940 should have worked itself out by 2000. However it is desirable to test this assumption. Bloom, Canning, and Sevilla (2004) examine a model where the error term is autocorrelated over time and has the structure

$$\varepsilon_{it} = (1 - \lambda)\varepsilon_{i,t-1} + \upsilon_{it} \qquad \textbf{(9)}$$

where υ_{it} is are independent identically distributed shocks, uncorrelated with $\varepsilon_{i,t-1}$. This implies that in the growth equation the residual is given by

$$\Delta\varepsilon_{it} = -\lambda\varepsilon_{i,t-1} + \upsilon_{it} \qquad \textbf{(10)}$$

And we can derive the growth model between two periods, 0 and 1, from equation (1) as

$$\Delta y_i = \delta + \pi\Delta x_i + \beta\Delta z_{it} + \lambda(u_i + \delta_0 + \pi x_{i0} + \beta z_{i0} - y_{i0}) + \upsilon_{i1} \qquad \textbf{(11)}$$

by replacing the lagged error term with the disequilibrium in the previous period. Note that in this case not all of the disequilibrium in one period is undone in the next period. Instead we have an error correction framework where income adjusts slowly to any disequilibrium between income and its steady-state level given initial health. While the residuals ε_{it} are correlated over

14

time, the shock υ_{it} is uncorrelated with any variables measured at time t-1, and these beginning-of-period variables, such as the predicted mortality change that depends only on mortality in the initial period, can be used as exogenous variables or instruments when estimating equation (11).

The model (11) is an encompassing model. It contains as a special case the models with uncorrelated errors in the level equation (1), which corresponds to $\lambda = 1$, giving equation (7). It also contains as a special case the model where the errors in the level relationship are a random walk, which corresponds to $\lambda = 0$ in equation (9). In this case the model reduces to the regression estimated by Acemoglu and Johnson since, the initial levels of health and income do not affect economic growth.

Over short time intervals we expect λ to be small, but over long time intervals λ will approach 1 as all of the initial disequilibrium is corrected. Cate (2004) provides a detailed discussion of the relationship between an underlying continuous time error correction mechanism and the value of λ when the process is measured at different discrete time intervals.

4. Regression Results

Table 2 reports regression results using the 47 country Acemoglu and Johnson data. For every set of results, the dependent variable is growth in per capita income during 1940-2000. In column 1 of Table 2 we replicate the Acemoglu and Johnson results from a regression of growth in income per capita on the change in log life expectancy. The point estimate of -1.142 (statistically significant) implies that a 10% increase in life expectancy leads to an 11.4% decrease in GDP per capita. In column 2 of Table 2 we follow Acemoglu and Johnson and instrument the change in log life expectancy with predicted mortality change. We find exactly the same result as Acemoglu and Johnson: the effect of growth in life expectancy on income growth is negative (a coefficient of -1.51), and statistically significant. With only two time periods the Acemoglu and Johnson fixed effects model in levels is equivalent to our first differenced model.

We begin by investigating the issue of omitted variable bias and test the robustness of the model to the inclusion of other possible explanatory variables. A key issue is whether initial health, as

15

measured by initial life expectancy, should be included in the growth model, as suggested by Figure 1. In column 3 of Table 2, we report the results of using OLS to estimate the parameters of a growth model that is augmented to include initial log life expectancy[8]. The inclusion of initial life expectancy reverses the result found in column 1: both initial log life expectancy and the change in log life expectancy are estimated to have positive and significant effects on economic growth. This means that the economies that grew fastest during 1940-2000 are those that had the highest life expectancies in 1940 and those that saw the fastest growth in life expectancy during 1940-2000.

We repeat the regression in column 3 but instrument growth in life expectancy with predicted mortality change and report the results in column 4 of Table 2. None of the resulting coefficient estimates are statistically significant. The problem with this specification is that once we condition on initial life expectancy, the instrument completely loses predictive power for changes in life expectancy in the first-stage regression. The Cragg-Donald F-Statistic of 0.137 implies that predicted mortality has virtually zero predictive power on subsequent changes in life expectancy when we control for initial life expectancy. Predicted mortality in 1940 is highly correlated with life expectancy in 1940 and once we condition on initial life expectancy the instrument becomes extremely weak.

In column 5 we add initial log income per capita to the model to estimate a more standard conditional convergence model. The addition of initial income has little impact on the OLS results. While we find a negative coefficient on initial GDP as is usual in conditional convergence models, the inclusion of initial GDP increases the estimated effect of both initial health and health improvements on economic growth. The fastest growing countries during 1940-2000 are those with low initial income, high initial life expectancy, and large increases in life expectancy during 1940-2000. We also try to estimate the parameters of this model using change in predicted mortality as an instrument for the improvement in life expectancy; the results are reported in column 6 of Table 2. Once again, the instrument is too weak in the first stage to allow for a proper identification of the model.

---

[8] This is equivalent to adding initial life expectancy interacted with a year dummy as an explanatory variable in the levels equation (1). This interaction is the same as carried out by Acemoglu and Johnson when they consider the robustness of their model to the inclusion of other variables.

One could argue that the instrument is invalid in column 3 because it is correlated with an relevant omitted variable, initial log life expectancy. By this logic, the instrument is valid in columns 4 and 6 since we control for initial health. However the instrument is now too weak to identify the effect of growth in life expectancy since it is highly correlated with initial life expectancy. These results imply that the Acemoglu and Johnson result is not robust to the inclusion of additional explanatory variables that are plausibly related to economic growth. With the addition of initial health as an explanatory variable in the growth regression the predicted mortality instrument is too weak to give meaningful results.

An alternative, however, is to estimate the parameters of the relationship set out in equation (11). This contains initial income per capita as well as initial log life expectancy and the growth in life expectancy as explanatory factors for the growth in income per capita. However, this model also has a parameter restriction not employed in the estimates reported in columns (5) and (6) of Table 2. The coefficients on the growth in life expectancy and initial log life expectancy, in the error correction component, are the same. This parameter restriction aids identification of the model.

Table 3 reports nonlinear least squares estimates of the parameters of equation (11) on the same data as used in Table 2 to explain growth in income per capita during 1940-2000. The effect, π, of initial life expectancy and the growth in life expectancy on economic growth are the same. We estimate a large effect of health on economic growth. In column 1 of Table 3 we estimate the relationship with non-linear least squares. We find a value of λ that is not significantly different from unity, indicating that all of the disequilibrium in 1940 between income and health appears to be corrected by the year 2000 through adjustments in income. This means that the income shock in 1940 and the shock in 2000 appear to be uncorrelated. In addition, the coefficient λ appears to be significantly different from zero, which is the condition Acemoglu and Johnson require for their approach to be valid (so that initial health and income can be excluded from the growth regression). We find that the coefficient on health (the common coefficient on initial health and the change in health) is positive, indicating a positive correlation between health and economic growth.

The main concern here, however, is that the change in log life expectancy Δx_i in equation (11) is endogenous to the growth in income. In column 2 of Table 3 we therefore instrument the change in log life expectancy with the predicted change in mortality. Note that this is now a plausibly valid instrument because the shock to the growth equation (11) depends only on the income shock in 2000, while the change in predicted mortality depends only on disease specific mortality measured in 1940. Using this instrument we again find positive and significant effects of health on growth. In columns 4 and 6 of Table 2 we estimate the effects of initial health and the change in health during 1940-2000 separately and the instrument is very weak. In column 2 of Table 3 identification essentially comes from the equality of the parameters on initial health and the growth in health in equation (11).

In the regression in column 2 of Table 3 we treat initial life expectancy and initial income per capita in 1940 as exogenous, which means that they are assumed to be uncorrelated with the income shock υ_{i1} in 2000. Rather than using predicted mortality change as our instrument for change in life expectancy we can simply use initial life expectancy as the instrument. While initial life expectancy is in the regression already, this instrument leads to identification because of the parameter restriction in equation (11). The results from using initial life expectancy as the instrument, reported in column 3 of Table 3, are very similar to those found by using predicted mortality change. In addition to the change in life expectancy being endogenous, we can also consider initial life expectancy as endogenous and instrument both with the change in predicted mortality. This corresponds to the argument that both initial life expectancy and the growth in life expectancy are potentially endogenous, but the change in predicted mortality is a valid instrument. The results of this regression are reported in column 4 of Table 3, and as before, the effect of health on income is positive and significant.

The results in Table 3 suggest that the coefficient λ on the error correction component of growth is unity. This means that the error terms in levels in equation (1) for 1940 and 2000 appear to be uncorrelated. In each regression in Table 3 we find a coefficient that is not significantly different from one. In Table 4 we repeat the regressions in Table 3 but impose the condition that $\lambda = 1$. The resulting estimated parameter for the effect of health on income is positive and significant in

18

each specification. Column 4 of Table 4 corresponds most closely to Acemoglu and Johnson's original approach. We substitute out the 1940 income shock to give equation (7), where we argue the predicted mortality change instrument becomes valid. We then instrument both initial log life expectancy and the change in log life expectancy with the predicted mortality instrument. The result is a large, significant, positive coefficient for the effect of life expectancy on income.

The estimate of the health effect is very sensitive to the extent of error correction. For low values of λ there is little effect of initial health on economic growth. In this case the negative relationship between the growth in life expectancy and the growth in income per capita, shown in Figure 2, dominates the estimate, and health improvements appear to lower income per capita. For large values of λ initial health has a greater impact on income growth, and the positive association between initial health and subsequent economic growth, shown in figure one, dominates the estimate and health improvements appear to raise income per capita. The critical value for λ is around 0.75, with larger values being required to generate positive effects of health improvements on income. We estimate a value of λ close to one, but the sensitivity of the estimated health coefficient to the value of λ should give us some pause. The fundamental reason for this sensitivity is the high degree of co-linearity between initial life expectancy, and the growth in life expectancy, shown in figure 2. It is intrinsically difficult to estimate separately the effects of the initial level, and the growth rate, of life expectancy using this data.

The results in Tables 2-4 require a large number of assumptions and are sensitive to the specification used; we do not regard them as definitive. Our goal is only to put the Acemoglu and Johnson results in context, and show they can be reversed in a more general model. Further generalizations are possible. A richer error structure than the simple autoregressive process set out in equation (9) would lead to a more complex model. In estimating equation (11) we neglect the fixed effects and the effect of other explanatory variables. It is difficult to include these fixed effects explicitly. In the long difference model we have only one observation in first differences. In a panel data setting with shorter time intervals and more observations we have dynamic panel with a lagged endogenous variable, leading to well-known difficulties (Nickell 1981; Arellano

and Bond 1991; Blundell and Bond 1998)[9]. Bloom, Canning, and Sevilla (2004) estimate the model set out in equation (11) in a panel of countries using a set of factors that are fixed within countries over time, such a geography and climate, to capture country-specific effects, and include other time-varying factors that may affect income per capita, such as capital accumulation. They continue to find a positive effect of health on economic growth, suggesting that the positive results found here are robust to including these variables.

5. Conclusion

The Acemoglu and Johnson (2007) results are striking. The idea that health improvements actually reduce income per capita is counterintuitive and, if true, of great practical importance with respect to the allocation of resources to the promotion and protection of population health. We argue here that the Acemoglu and Johnson result hinges critically on a specification in which economic growth depends on contemporaneous health improvements only, so that lagged health has no effect. There are strong reasons to expect that the response of income to health improvements is not instantaneous, and that there is a lagged effect. Allowing for this lagged effect produces positive estimates for the effects of both lagged health and of improvements in health on the growth of income per capita, and thus reverses the results found by Acemoglu and Johnson.

While we find positive results, we are left with a situation in which we are not yet able to make definitive inferences about the effect of population health on economic growth. There is a strong positive correlation between a country's initial life expectancy and its subsequent economic growth, but it is difficult to offer compelling evidence that the effect is causal. We still lack a randomized intervention, or an indisputably compelling instrument, for initial health at the macroeconomic level that would allow us to identify this effect of initial life expectancy as causal and not the result of omitted variable bias.

---

[9] Method of moments give asymptotically unbiased estimators in this setting but we found that a Monte Carlo study using generated data gave very unreliable results.

While we disagree with the results found by Acemoglu and Johnson we think their paper has identified three crucial issues that are not yet resolved. The first issue is the endogenity of health and the need for modeling that allows causal inference. The second issue is timing. Any effect of health improvements that increases income levels clearly takes some time - with slow convergence to a steady state, rather than a contemporaneous jump, the best we can hope for. The third is the indirect effect of health on population size, and the resulting pressure on scarce resources and income levels. Previous studies may have erred in controlling for population effects rather than treating these as endogenous.

We have focused on the overall effect of health on income, which includes labor productivity and savings, as well as population effects. We agree that the population mechanism may lead to a negative effect of health on income in the short run, particularly since the labor productivity effects of early childhood health may take years to realize. However, in the longer run, improvements in mortality tend to lead to reduced fertility. A central question is whether induced reductions in fertility are sufficient to offset the population increase due to lower mortality, and how quickly the fertility response materializes (Ashraf, Lester et al. 2008).

.

TABLE 1: CORRELATION MATRIX ACEMOGLU AND JOHNSON SAMPLE OF 47 COUNTRIES

	Growth in income per capita 1940-2000	Growth in life expectancy 1940-2000	Log life expectancy 1940	Log GDP per capita 1940	Predicted mortality 1940
Growth in income per capita 1940-2000	1				
Growth in life expectancy 1940-2000	-0.44***	1			
Log life expectancy 1940	0.50***	-0.97***	1		
Log GDP per capita 1940	0.13	-0.76***	0.81***	1	
Predicted mortality 1940	-0.34**	0.69***	-0.70***	-0.63***	1

NOTE.—Table 1 shows pairwise correlation coefficients for the variables listed. *** p<0.01, ** p<0.05, * p<0.10. Based on the data for the 47 countries used in Acemoglu and Johnson (2007).

TABLE 2: GROWTH IN LIFE EXPECTANCY AND INCOME PER CAPITA, ADDING INITIAL CONDITIONS

	GROWTH IN INCOME PER CAPITA, 1940-2000					
	(1)	(2)	(3)	(4)	(5)	(6)
Estimator	OLS	IV	OLS	IV	OLS	IV
Growth in life expectancy 1940-2000	-1.142***	-1.506***	3.684***	-21.56	5.774***	-21.44
	(0.27)	(0.40)	(1.29)	(81.3)	(1.18)	(66.8)
Log life expectancy 1940			3.769***	-15.23	6.848***	-15.18
			(0.94)	(61.4)	(1.00)	(54.3)
Log GDP per capita 1940					-0.689***	0.0146
					(0.11)	(1.81)
Observations	47	47	47	47	47	47
R-squared	0.19	0.17	0.32	-3.13	0.60	-3.09
Cragg-Donald F-Statistic		60.84		0.136		0.211
Critical value for F-statistic		16.38		16.38		16.38

NOTE.─Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. In IV estimates change in log life expectancy is instrumented with predicted mortality change. Based on the Acemoglu and Johnson sample of 47 countries.

TABLE 3: GROWTH IN LIFE EXPECTANCY AND INCOME PER CAPITA, ERROR CORRECTION MODEL

		GROWTH IN INCOME PER CAPITA, 1940-2000			
		(1)	(2)	(3)	(4)
Estimator		NLS	IV	IV	IV
Intercept	α	-33.39**	-53.10**	-55.86*	-41.91**
		(3.385)	(13.31)	(15.08)	(10.71)
Effect of health on income	π	9.143**	14.96**	15.78**	11.24**
		(1.067)	(3.943)	(4.490)	(3.353)
Convergence term	λ	1.115**	0.959**	0.947**	1.076**
		(0.071)	(0.065)	(0.065)	(0.135)
R^2		0.560	0.494	0.473	0.584
Endogenous			Growth in life expectancy	Growth in life expectancy	Growth in life expectancy, Initial log life expectancy
Instrument			Change in predicted mortality	Initial log life expectancy	Change in predicted mortality

NOTE.─Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Based on the Acemoglu and Johnson sample of 47 countries. Estimated specifications are based on equation (11).

TABLE 4: GROWTH IN LIFE EXPECTANCY AND INCOME PER CAPITA, ERROR CORRECTION MODEL

		GROWTH IN INCOME PER CAPITA, 1940-2000			
Estimator		(1) NLS	(2) IV	(3) IV	(4) IV
Intercept	α	-39.37** (2.951)	-46.01** (4.395)	-46.21** 4.39532	-51.37** 5.81075
Effect of Health on Income	π	11.28** (0.685)	12.83** (1.018)	12.87** (1.018)	14.08** (1.348)
Convergence term	λ	1	1	1	1
R^2		0.496	0.532	0.532	0.541
Endogenous			Growth in life expectancy	Growth in life expectancy	Growth in life expectancy, Initial log life expectancy
Instrument			Change in predicted mortality	Initial log life expectancy	Change in predicted mortality

NOTE.─ Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Based on the Acemoglu and Johnson sample of 47 countries. See equation (11) for the specification being estimated. We impose the condition $\lambda = 1$.

FIGURE 1: LIFE EXPECTANCY 1940 AND GROWTH IN INCOME PER CAPITA 1940-2000
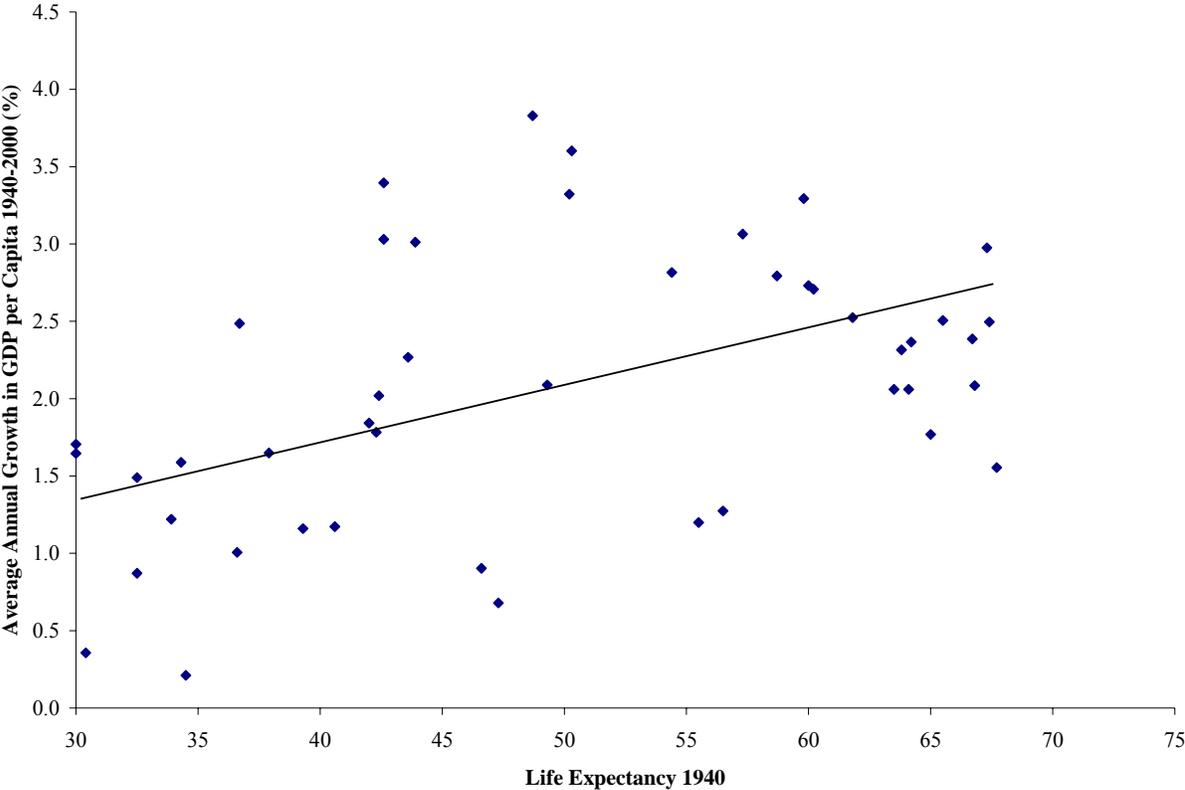
FIG. 1.— Life expectancy 1940 and average annual growth in income per capita 1940-2000 in the Acemoglu and Johnson baseline sample of 47 countries.

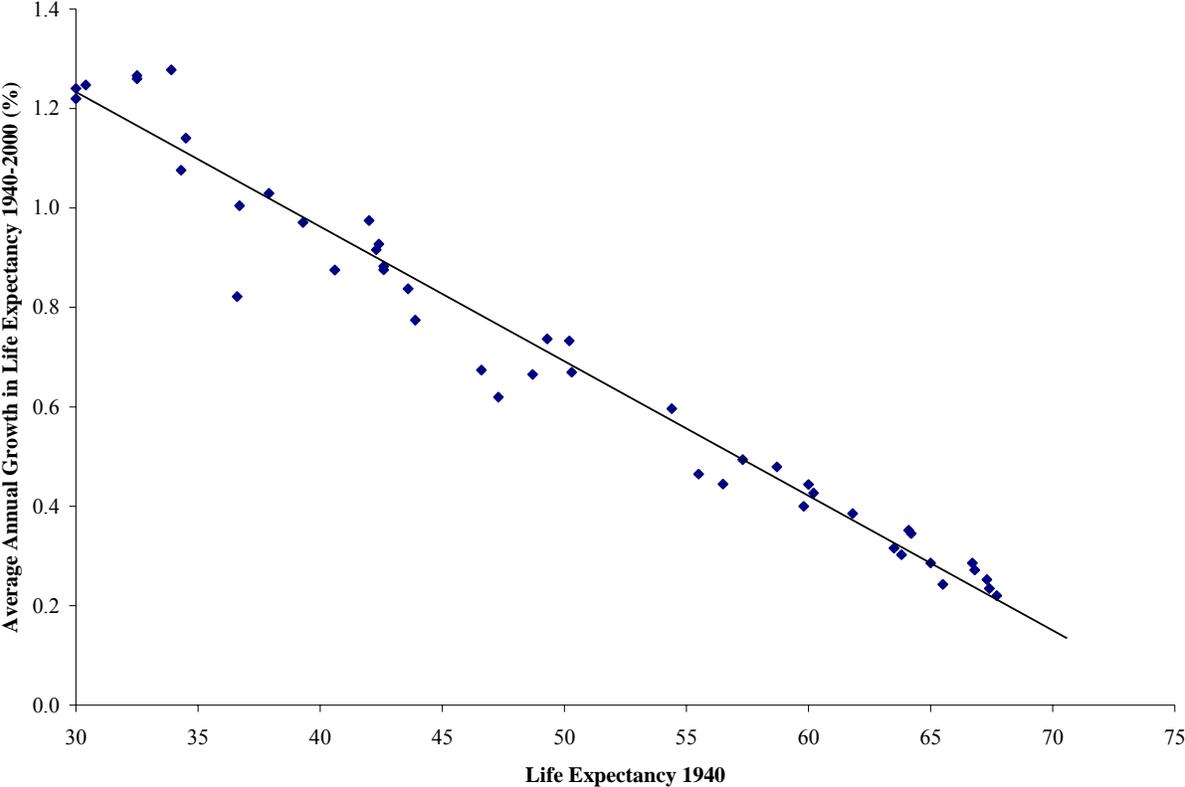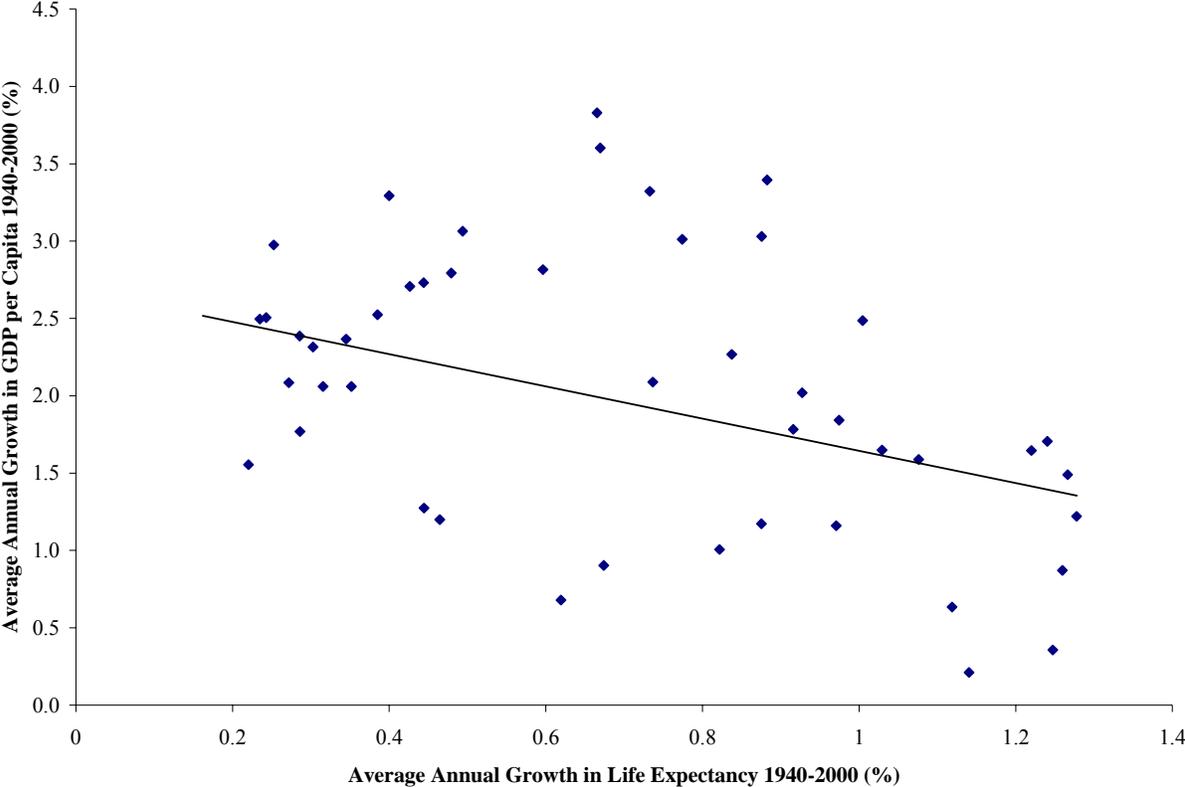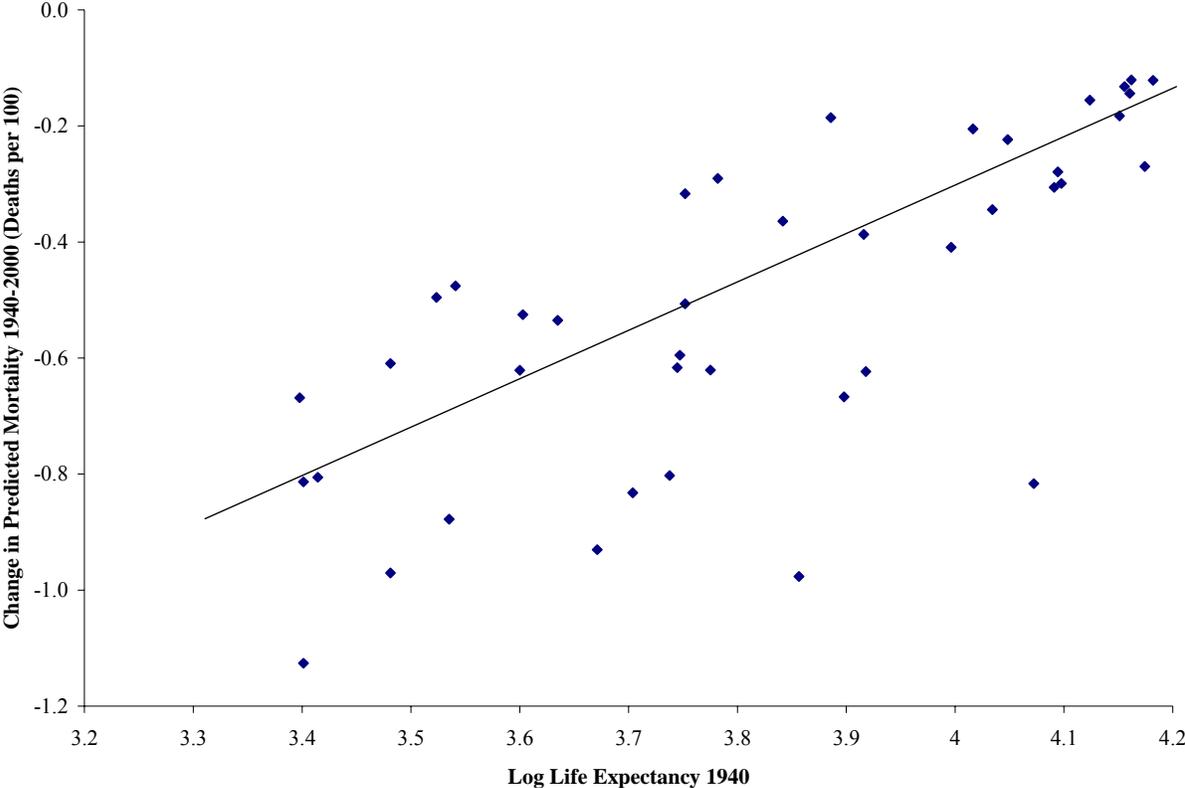**FIGURE 2: LIFE EXPECTANCY 1940 AND GROWTH IN LIFE EXPECTANCY 1940-2000**

FIG. 2.— Average annual growth in life expectancy 1940-2000 and life expectancy 1940 in the Acemoglu and Johnson baseline sample of 47 countries.

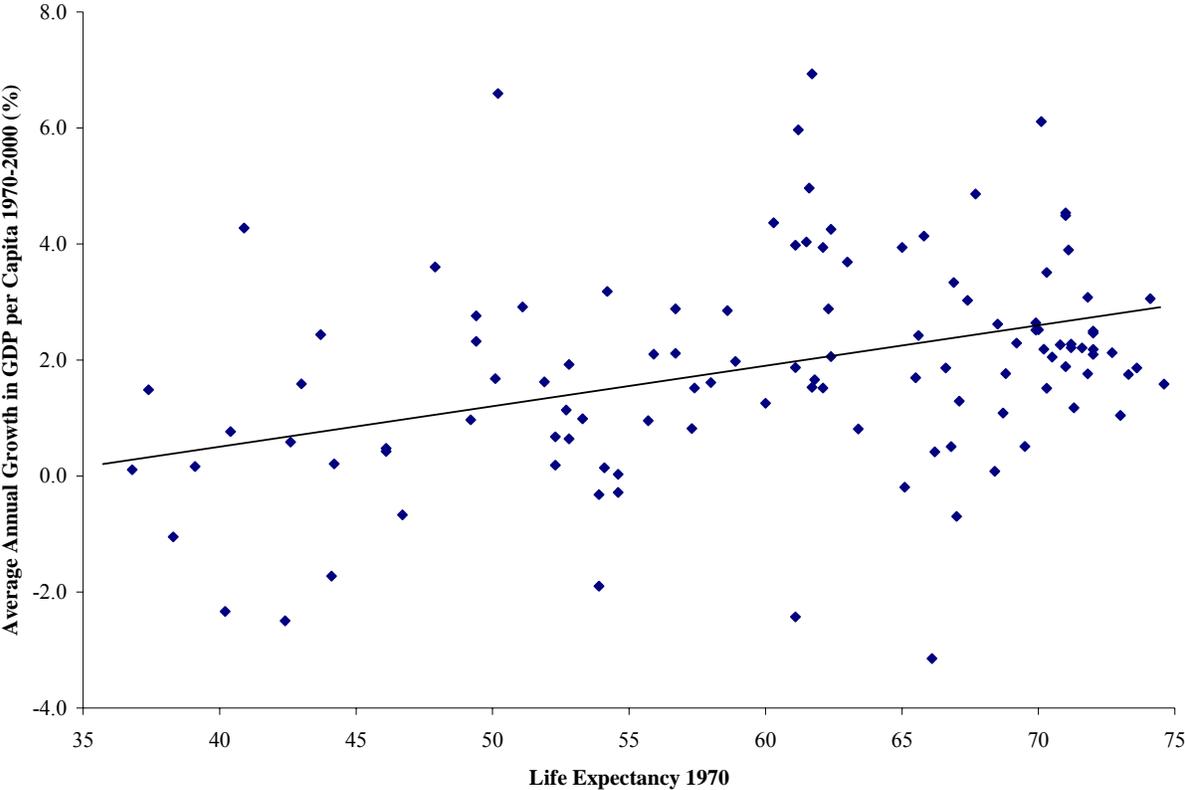**FIGURE 3: GROWTH IN LIFE EXPECTANCY AND GROWTH IN INCOME PER CAPITA 1940-2000**

FIG. 3.─ Average annual growth in life expectancy and GDP per capita 1940-2000 in the Acemoglu and Johnson baseline sample of 47 countries.

FIGURE 4: LOG LIFE EXPECTANCY 1940 AND CHANGE IN PREDICTED MORTALITY 1940-2000

FIG. 7.─ Absolute change in predicted mortality 1940-2000 and log life expectancy 1940 in the Acemoglu and Johnson baseline sample of 47 countries.

FIGURE 4: LIFE EXPECTANCY 1970 AND GROWTH IN INCOME PER CAPITA 1970-2000

FIG. 4.─ Life expectancy 1970 and average annual growth in income per capita 1970-2000 in the extended sample of 112 countries. Source: GDP per capita: Penn World Tables 6.2; life expectancy: WDI 2007.
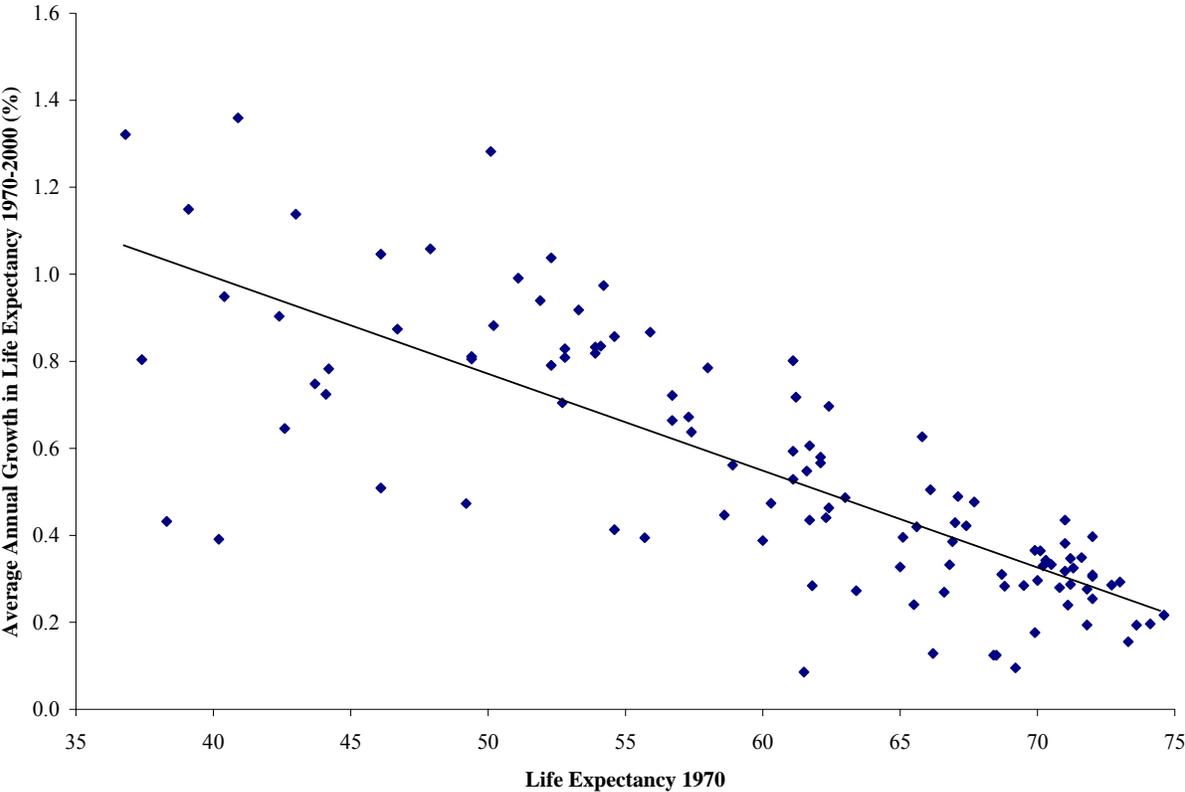
FIG. 5.─ Average annual growth in life expectancy 1970-2000 and life expectancy 1970 in the extended sample of 112 countries. Source: GDP per capita: Penn World Tables 6.2; life expectancy: WDI 2007.

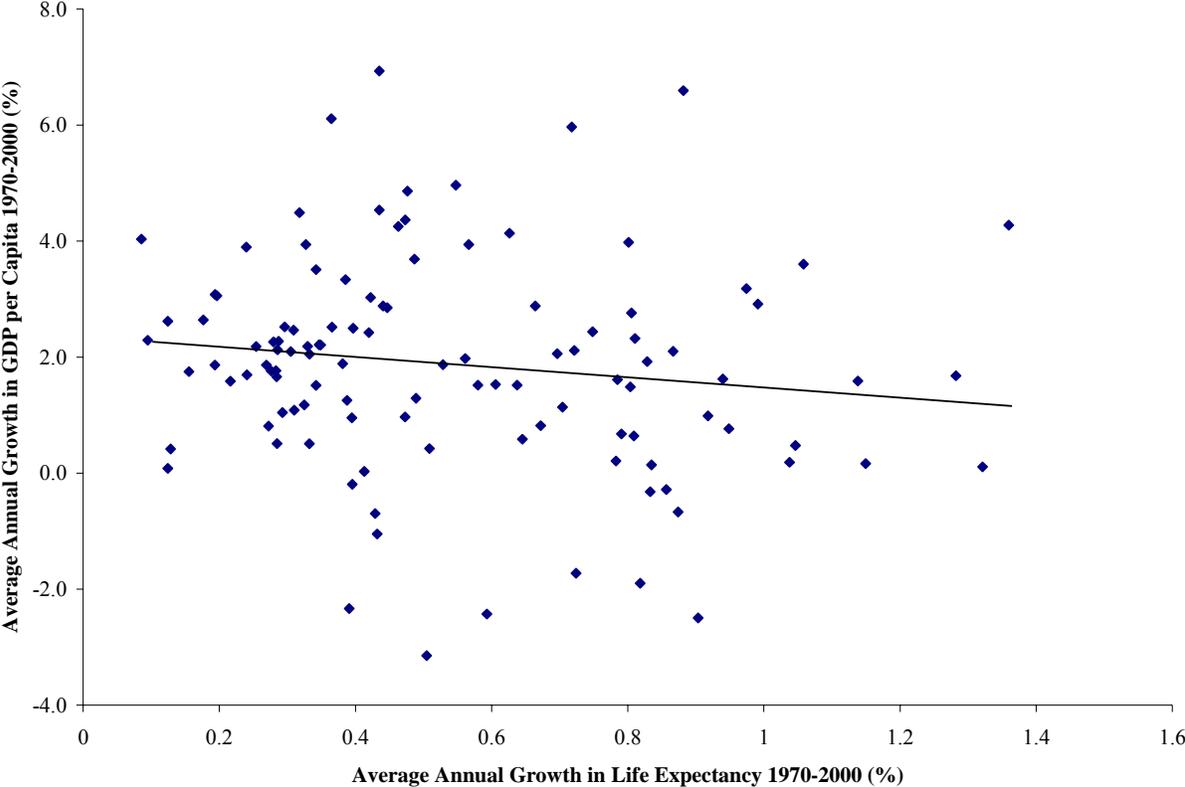**FIGURE 6: GROWTH IN LIFE EXPECTANCY AND GROWTH IN INCOME PER CAPITA 1970-2000**

FIG. 6.─ Average annual growth in life expectancy and GDP per capita 1970-2000 in the extended sample of 112 countries. Source: GDP per capita: Penn World Tables 6.2; life expectancy: WDI 2007.

References

Acemoglu, Daron and Simon Johnson (2007). "Disease and Development: The effect of life expectancy on economic growth." *Journal of Political Economy* 115(6): 925-985.

Angrist, Joshua D. and Jörn-Steffen Pischke (2009). *Mostly Harmless Econometrics*. Princeton: Princeton University Press.

Arellano, Manuel and Stephen Bond (1991). "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations." *Review of Economic Studies* 58(2): 277-97.

Ashraf, Quamrul H., Ashley Lester and David N. Weil (2008). "When Does Improving Health Raise GDP " *NBER Macreconomics Annual* 23: 157-204.

Barker, David J.P. (1992). *The Fetal and Infant Origins of Adult Disease*. London: BMJ Books.

Barro, Robert J. and Xavier Sala-i-Martin (2004). *Economic Growth*. Cambridge, MA: MIT Press.

Behrman, Jere R. and Mark R. Rosenzweig (2004). "The Returns to Birthweight." *Review of Economics and Statistics* 86(2): 586-601.

Bleakley, Hoyt (2003). "Disease and Development: Evidence from the American South." *Journal of the European Economic Association* 1(2-3): 376-386.

Bleakley, Hoyt (2006). "Malaria in the Americas: A Retrospective Analysis of Childhood Exposure." *Mimeo University of Chicago*.

Bleakley, Hoyt (2008). "Comment on "When Does Improving Health Raise GDP"." *NBER Macreconomics Annual* 28(February): 205-220.

Bloom, David E., David Canning and Jaypee Sevilla (2004). "The Effect of Health on Economic Growth: A Production Function Approach." *World Development* 32(1): 1-13.

Blundell, Richard and Stephen Bond (1998). "Initial Conditions and Moment Restrictions in Dynamic Panel Data Models." *Journal of Econometrics* 87(1): 115-43.

Case, Anne, Angela Fertig and Christina Paxson (2005). "The Lasting Impact of Childhood Health and Circumstance." *Journal of Health Economics* 24(2): 365-389.

Cate, Arie ten (2004). "Refinement of the Partial Adjustment Model Using Continuous-time Econometrics." *CPB Discussion Papers* 41.

Crimmins, Eileen M. and Caleb E. Finch (2006). "Infection, Inflammation, Height, and Longevity." *Proceedings of the National Academy of Sciences* 103(2): 498-503.

Cutler, David, Angus Deaton and Adriana Lleras-Muney (2006). "The Determinants of Mortality." *Journal of Economic Perspectives* 20(3): 97-120.

Durlauf, Steven, Paul Johnson and Jonathan Temple (2005). Growth Econometrics. Handbook of Economic Growth. P. Aghion and S. Durlauf. Amsterdam, Elsevier.

Heckman, James J. (2007). "The economics, technology, and neuroscience of human capability formation." *Proceedings of the National Academy of Sciences of the United States of America* 104(33): 13250-13255.

Heston, Alan, Robert Summers and Bettina Aten (2006). "Penn World Table Version 6.2." *Center for International Comparisons of Production, Income and Prices at the University of Pennsylvania*.

Hoddinott, J., J. Maluccio, et al. (2008). "Effect of a nutrition intervention during early childhood on economic productivity in Guatemalan adults." *Lancet* 371(9610): 411-416.

Kremer, Michael and Edward Miguel (2004). "Worms: Identifying Impacts on Education and Health in the Presence of Treatment Externalities " *Econometrica* 72(1): 159 - 217.

Nickell, Stephen (1981). "Biases in Dynamic Models with Fixed Effects." *Econometrica* 49(6): 1417-1426.

Packard, Randall M. (2007). *The Making of a Tropical Disease: A Short History of Malaria*. Baltimore, Maryland: The Johns Hopkins University Press.

Schultz, Paul T. (2005). Productive Benefits of Health. Health and Economic Growth. G. López-Casasnovas, B. Rivera and L. Currais. Cambridge MA, MIT Press.

Snowden, Frank M. (2006). *The Conquest of Malaria: Italy, 1900-1962*. New Haven, Connecticut: Yale University Press.

Weil, David (2007). "Accounting for the Effect of Health on Economic Growth." *Quarterly Journal of Economics* 122(3): 1265-1306.

World Bank (2007). World Bank Development Indicators CD-ROM.

Appendix I

By definition,

$$Cov(M^I_{it}, \varepsilon_{it}) = E\left[(M^I_{it} - E(M^I_{it}))(\varepsilon_{it} - E(\varepsilon_{it}))\right] \tag{12}$$

Now assuming that $E(\varepsilon_{it}) = 0$ we have

$$Cov(M^I_{it}, \varepsilon_{it}) = E(M^I_{it}\varepsilon_{it}) - E(M^I_{it})E(\varepsilon_{it}) = E(M^I_{it}\varepsilon_{it}) \tag{13}$$

Since we have only two periods, we can express the covariance term as

$$E(M^I_{it}\varepsilon_{it}) = E(M^I_{i0}\varepsilon_{i0})P(t=0) + E(M^I_{i1}\varepsilon_{i1})P(t=1) = E(M^I_{i0}\varepsilon_{i0})P(t=0), \tag{14}$$

since by construction predicted mortality at time 1 is zero so $M^I_{i1} = 0$. Hence, given that some observations in the dataset are at time zero (so that $P(t=0) > 0$), we have

$$Cov(M^I_{it}, \varepsilon_{it}) = 0 \iff E(M^I_{it}\varepsilon_{it}) = 0 \iff E(M^I_{i0}\varepsilon_{i0}) = 0 \iff Cov(M^I_{i0}, \varepsilon_{i0}) = 0 \tag{15}$$