Disease and Development: Comments on Acemoglu and Johnson (2006)*

Hoyt Bleakley†

July 16, 2006

Acemoglu and Johnson (2006) have produced an important and fascinating study that seeks to quantify the contribution of health to economic development. This question is of tremendous importance both for scientific and policy-related motivations. There remains significant doubt about the fundamental determinants of why some countries are rich and some are poor, and even more doubt about what can be done about it.

Deficiencies in health (and health policy) are recurrent hypotheses put forward to explain income differences across countries, research into the role of health is needed. However, the existing literature does not address this question satisfactorily. Macroeconomic studies of a cross section of countries or of a time series within a particular country simply do not measure the causal effect of health. Third factors affect both health and productivity, and income is likely a significant determinant of investments in health. Without cutting through this Gordian knot of identification, the extant literature is of limited usefulness.

Large-scale public-health interventions are a possible source of exogenous variation. Over the past two hundred years, such interventions have often received their initial impulse from advances in medical knowledge and were further catalyzed by funding from external entities (international philanthropies and donor-country governments). To some degree, these features alleviate concerns about reverse causality (from income to health).

Acemoglu and Johnson consider advances in and diffusion of technology *circa* 1950 that brought down mortality rates in poorer countries. During this period, a variety of new

*Remarks delivered at the NBER Summer Institute on Economic Fluctuations and Growth, July 16, 2006.

†Assistant Professor of Economics, Graduate School of Business, University of Chicago, 5807 South Woodlawn Avenue, Chicago, IL, 60637. Telephone: (773) 834-2192. Electronic mail: bleakley[at]chicagogsb[dot]edu

medicines (antibiotics, e.g.) and new techniques (residual-action spraying against mosquitos, for instance) were brought to bear on the health problems worldwide.

Take-up of new technology is, of course, endogenous, so the authors focus on countries that stood to benefit the most. (This is therefore an "intention to treat" design.) Operationally, this means that the countries with a higher burden of a particular disease received a more favorable shock when the technology for fighting that disease advanced.[1]

A simple summary of their study is as follows:

1. Technological progress decreased LDC mortality (*c.* 1950)
2. Life expectancy ↑, big time.
3. Population ↑.
4. GDP ↑, but by less than population, so...
5. GDP per capita ↓

The econometric analysis is careful and thorough. As a historical evaluation of this episode, these facts seem to be a good description of the data (although I flag a few unusual features of the results below). But for the past to be a guide for today's policy decisions, we need unpackage this result to understand this seeming disconnect between health and productivity.

The authors' interpretation of the GDP result is that other factors did not adjust when population grew. In terms of our usual symbols, $L \uparrow$, but

$$(K/L) \downarrow \quad \Rightarrow \quad (Y/L) \downarrow$$

where K might include physical capital or improved land. Note that this is a general-equilibrium effect. It stands in stark contrast with the results of many micro- and cohort-level studies that estimate large, positive effects of health on productivity, although presumably of a partial-equilibrium nature. If these latter effects were present, they were, by the authors' logic, more than fully offset by the "crowding" effect of population growth.

There is also a kind of uber-interpretation about the irrelevance of health in raising productivity to any significant degree in LDCs. Acemoglu and Johnson are careful about not saying this too strongly, but invariably some will come away with this conclusion.

[1]Because I use this methodology in my own work, selection effects militate against my being its best critic. The main issue is to check that pre-innovation disease rates are not correlated with pre-existing trends, which the authors do ably.

Are these the right interpretations of the results? Their hypothesis of fixed factors may be right for this episode, although I point out some limitations of this view below. As to the larger question, I argue that their particular slice of the data provides an incomplete picture of the role of health and income. My comments revolve around several themes, with my apologies to the reader for the nonlinearity of the organization to follow.

- Partial versus general equilibrium. (Certainly capital and land might be sluggish and/or inelastic. But are the estimated responses plausible under this interpretation?)
- The incidence of mortality across the life cycle. (Their instrument disproportionately affects early-life mortality. What effects might this have relative to adult and/or late-life mortality?)
- Mortality versus morbidity. (Their instrument is weighted towards mortality. How might this matter?)
- Cohort versus time. (Which is the right unit of observation for measuring the effects of health?)

An important assumption of their study is that of a single-index model of health. In other words, increases in life expectancy at one age are for practical purposes the same as increases in life expectancy at other ages, and that these changes have similar effects as any other improvement in health. These assertions are unlikely to be true, for both theoretical and empirical reasons. Therefore it is worth considering how to interpret their results without the single-index assumption.

How should life expectancy affect productivity? Gains in life expectancy at working ages increase the time during which childhood human-capital investments can be utilized. This constitutes the most most obvious (to me anyway) link between mortality and income.[2] Reductions in early-life mortality, on the other hand, do not affect this calculation. Indeed, the most direct effect is to reduce the cost to parents of having more descendents (the price of "quantity" in the usual lingo).[3] Absent any change in the price of "quality", we might reasonably expect a movement along the Beckerian quantity/quality tradeoff. In other words, standard price-theoretic arguments predict more kids, albeit with less average education.

[2]Note however that Hazan (2006) presents evidence that increases in (adult) life expectancy in the U.S. were not accompanied by increases in lifetime hours.

[3]See David Meltzer's 1992 dissertation for more on this point. See also Kalemli-Ozcan (2002), who argues that the second moment rather than the expected length of life could counteract this quantity-price effect.

The innovations studied by Acemoglu and Johnson largely improved mortality in infancy and early childhood.[4] Given this fact, these early-life increases in life expectancy induce a reaction in fertility that is consistent the argument above, as shown in Figure 1, Panel A. While there's not enough good data on education to measure the quality response, a quantity response is evident: contemporaneous with the mortality decline, birth rates are higher, and the (log) number of births remains higher for the 40-year horizon of their estimates.[5]

As a point of comparison, consider the optimal response to an intervention in health that reduces childhood morbidity with little to no effect on mortality. The price of quantity is unaffected, but the price of quality drops, if childhood sickness is an impediment to learning. We would expect fertility to decline in this case as parents move along the quantity/quality tradeoff. In joint work with Fabian Lange (Bleakley and Lange, 2006), we test precisely this hypothesis by considering the Rockefeller-funded campaign against hookworm in the American South, *circa* 1913. In that episode, hookworm was almost exclusively a childhood disease, and had neglible effects on mortality. We show that areas with high hookworm infection rate prior to eradication also had higher fertility, but these same areas saw reductions in birth rates following the anti-hookworm campaign. This result can be seen in Figure 1, Panel B.

I now turn to the question of cohort versus time. The benefits of the mortality decline were concentrated among the young for two reasons. First, much of gains in life expectancy occurred at younger ages. Second, because most of human-capital investments are made in childhood, most likely only the young could react appreciably to any gains in adult life expectancy. Both of these mechanisms suggest a cohort rather than a time effect, and highlight the importance of childhood exposure to the mortality reduction. Consider a campaign that starts in year zero and takes effect instantaneously. Cohorts born after this date will be exposed to the campaign for their entire childhood. On the other hand, those cohorts who were already adults in year zero will have no childhood exposure to the campaign,

[4]Tuberculosis is an exception, which tended to strike adults. Malaria mortality is also concentrated among infants and children, although the story is a bit more complicated. While some types of malaria have low case-fatality ratios, their instrument is based on reported deaths, and consequently is implicitly weighted towards malarious areas with high mortality (probably from *falciparum*) rather than high morbidity. I return to this point below.

[5]Ten to twenty years later, net birth rates essentially return to their earlier levels. Nevertheless, total births remain high because of the earlier expansion of population. Oddly, birth rates 30–40 years are substantially depressed, suggesting that people want to return to their 1940s population density (relative to other countries). One interpretation is that it takes agents 30+ years to figure out how to calibrate their fertility to obtain the optimal population density, although I am concerned that the result might reflect some data problem.

while the "in between" cohorts will be partially exposed during childhood, as shown in Figure 2, Panel A. Some time would have to pass before these healthier cohorts filled out the working-age population, and this transition path is shown in Panel B, which considers an intervention in the late 1950s (the beginning of the intensive years of the malaria-eradication campaigns in Latin America). As seen in the graph, only around 20% of the effect is realized by 1980, and the full gains are not seen until almost 2010, based on the timing assumptions above.

In my own work (Bleakley 2006a and 2006b), I show that this cohort model of childhood-exposure is a strong predictor of later-life productivity. Figure 3, which graphs income and pre-eradication malaria by birth cohort, shows a marked shift coincident with the campaign against malaria in the Southern United States in the 1920s. Similar results are found for the Rockefeller-funded campaign against hookworm a decade earlier. Moreover, some question the relevance of historical evidence from the U.S. with to developing countries today, so I repeat the cohort-based analysis of the 1950s malaria-eradication campaigns in Brazil, Colombia, and Mexico, and I find broadly similar results.

Having subjected the reader to all this shameless self promotion, I should point out that I doubt that confusing cohort with time does much to expain differences between our results. Consider their estimates of how GDP per capita responded over the forty years following the improvements in life expectancy. These coefficients are plotted in Figure 4. No up-tick is evident, even in the later years.

This brings us back to the author's favored interpretation for the results: that the crowding effect of additional population overwhelmed any partial-equilibrium gains in productivity. If others factors cannot adjust, this general-equilibrium effect of crowding is an inevitable result (albeit of uncertain magnitude). Their estimates for GDP are insigificantly different from zero. This is surprising: apparently these countries cannot do *anything* with that extra population. Bear in mind that we are measuring not labor income, but GDP, which should include income accruing to the other factors. In fact, if we had measures of the other factors' income, would we observe some twenty dollar bills (peso notes?) on sidewalk? (Indeed, an interesting plausibility check would be to predict the return to capital based on their calculation at the end of the paper.)

Capital being fixed in "medium run" is a pretty common assertion in micro classes (mine included), and so it gets no grief from me. But turn your eyes back to Figure 4. Their estimate for the contemporaneous response of GDP per capita to life expectancy is about -1. Looking at the response of GDP over the following four decades, we see that the response

is still -1, essentially the same as the immediate effect. Is forty years still the medium run? Probably not.

On the other hand, it is hardly controversial to say that land is supplied inelastically. Many of these countries were highly dependent on agriculture in the 1940s, but this changed markedly in subsequent decades. Latin America, for example, went from 40% urban at mid-century to almost 80% today. This is a potential margin of adjustment that could be investigated. In any event, it certainly takes us out of the realm of the stylized, land-dependent economy. Note, moreover, that the green revolution was raising the productivity of land during this period.

It bears mentioning that some of the existing "micro" literature does consider general-equilibrium effects. In a randomized evaluation of deworming treatments in Kenya, Miguel and Kremer (2004) find evidence of spillovers within schools from treated to untreated pupils, not all of which could be accounted for by disease transmission. (Although they do not examine income, they do have test scores.) My own work considers outcomes at the cohort level, so any general-equilibrium effect within the cohort is built into the estimate. Moreover, if the healthier cohorts simply displaced older workers (because of ranking or signalling), we would expect a different pattern to the coefficients in Figure 3. A pure ranking/displacement mechanism would generate income differences by childhood exposure for those born within a generation of the intervention. But this mechanism would not affect the endpoints of the time series. Why? Because people born in 1820 never worked in the same labor market with cohorts exposed to the anti-malaria campaign as children. Similarly, people born in 1960 always worked alongside those fully exposed to the treatment. As seen in the figure, the end points tend to be a bit closer together, but a gap remains. (Future work will test this more rigorously by decomposing these results into cohort $\times$ time and introducing own exposure and average exposure as separate regressors.)

Another key difference that emerges between these studies and Acemoglu and Johnson is along the morbidity/mortality dimension. The former set considers parasitic infections that tend to have low case-fatality rates. The latter is weighted towards more lethal diseases. In my own work, I have found an interesting natural experiment that bridges this gap. Within Colombia, most of the malarious areas were afflicted with *vivax* malaria, the high-morbidity strain. However, significant portions of the country suffered from elevated rates of *falciparum*, a malaria parasite associated with high mortality. When I estimate an interacted model, I find that eradicating *vivax* malaria produced substantial gains in human capital and income. On the other hand, estimates indicate no such gains from eradicating *falciparum*.

(See Bleakley, 2006b, for these estimates.)

Let me suggest an alternative synthesis of the story so far. Mortality has for a long time played a central role in human-capital theory (including numerous studies building on the seminal work of Becker and ben Porath). On the other hand, Acemoglu and Johnson (2006) argue that mortality has a limited effect on aggregate productivity. But childhood morbidity seems to make a difference.

Consider a bare-bones version of the FOC for human-capital choice:

$$c'(H) = \sum \beta^t f'_t(H)$$

where $c()$ is the cost of acquiring human capital (H), $f()$ is economic return to H, and β is the discount rate. Assume the usual things about $c()$, *i.e.*, $c'(H) > 0$, $c''(H) > 0$. Mortality is like the β term, death being the ultimate form of depreciation. Childhood morbidity affects $c(H)$ by impeding learning.[6]

This model is illustrated in Figure 5 with hypothetical functional forms. The top panel presents the familiar two-curves-crossing graph. The discounted sum of marginal benefits from human capital appears as a horizontal, dashed line. The marginal-cost function is the red line, and displays the property that acquiring above a certain level of human capital is effectively impossible. The middle panel considers the effect of decreasing working-age mortality. This decreases the discount rate applied to future returns to schooling, thereby shifting the marginal-benefits curve up. As shown in the figure, at some point additional increases in life expectancy yield minimal changes in human capital if morbidity causes the marginal-cost function to rise too rapidly. The bottom panel, in contrast, shows how a decrease in morbidity can increase human-capital investment, in effect relaxing the constraint shown in the middle panel.

In sum: this is a good paper, but we need to further open up the black box of health when thinking about the relationship between disease and development.

[6]This definition of cost includes opportunity cost, which is presumably affected by being sick as well. Thus I am implicitly assuming that childhood morbidity (anemia, in particular) has a greater effect on the ability to learn increasingly complicated concepts than it does on the return to child labor. Note further that I am not necessarily equating human capital with years of education. Increasing the rate of learning in formal schooling has ambiguous effects on the number of years one sits in school.

References

Acemoglu, Daron and Simon Johnson (2006). "Disease and Development: The Effect of Life Expectancy on Economic Growth." Working Paper 12269, National Bureau of Economic Research, Cambridge, Mass. June.

Bleakley, C. Hoyt (2007). "Disease and Development: Evidence from Hookworm Eradication in the American South." *The Quarterly Journal of Economics.* February **122** (1), 73–117.

Bleakley, Hoyt (2006). "Malaria in the Americas: A Retrospective Analysis of Childhood Exposure." March. Unpublished manuscript.

——— and Fabian Lange (2006). "Chronic Disease Burden and the Interaction of Education, Fertility and Growth." Working Paper 121, Bureau for Research and Economic Analysis of Development, Cambridge, Mass. May.

Hazan, Moshe (2006). "Longevity and Lifetime Labor Input: Data and Implications." June. Unpublished manuscript, accessed at http://www.nber.org/confer/2006/si2006/efjk/hazan.pdf.

Kalemli-Ozcan, Sebnem (2002). "Does Mortality Decline Promote Economic Growth?" *Journal of Economic Growth.* **7.**

Meltzer, David (1992). "Mortality decline, the Demographic Transition, and Economic Growth." Doctoral dissertation, University of Chicago December.

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

Figure 1: Fertility Responses to Mortality versus Morbidity

Panel A: Response to Mortality Decline, Zero to Forty Years Later

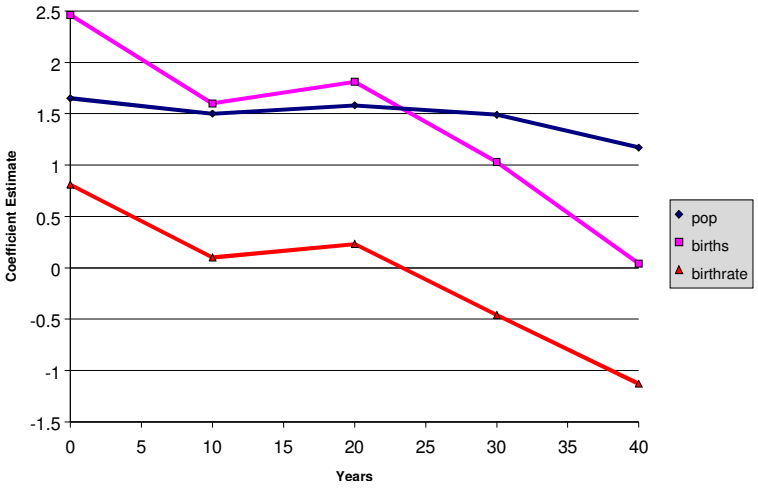

Panel B: Fertility Response to Hookworm Eradication *circa* 1913

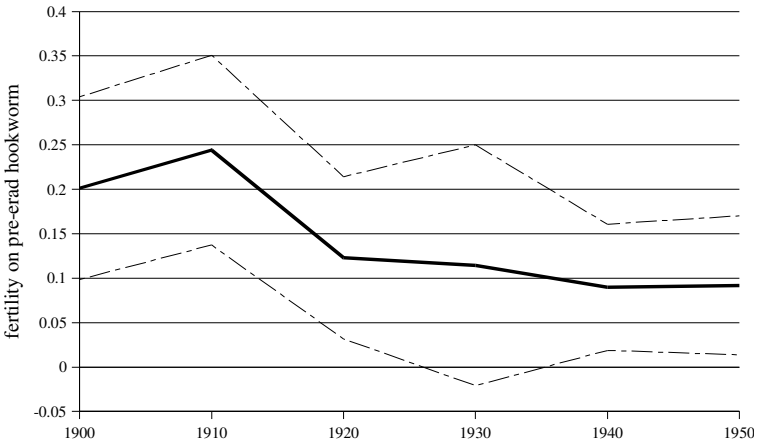

Notes: Panel A plots the 2SLS coefficients on life expectancy at various lags for the indicated outcomes. The estimates are drawn from various tables in Acemoglu and Johnson (2006). The birth rate is the (log) difference between the birth and population series. Panel B is drawn from Bleakley and Lange (2006). The y axis plots the year-specific coefficients of fertility on the *circa* 1913 hookworm-infection rate (solid line), plus the 95%-confidence intervals (dashed lines). The x axis is the Census year.

Figure 2: Childhood Exposure to Eradication Campaign

Panel A: Childhood Exposure by Birth Cohort

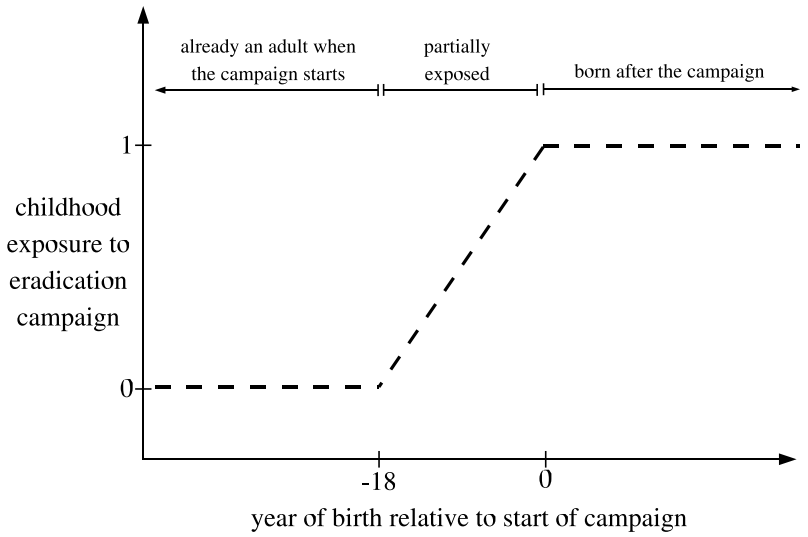

Panel B: Comparison of Cohort versus Time Effects:

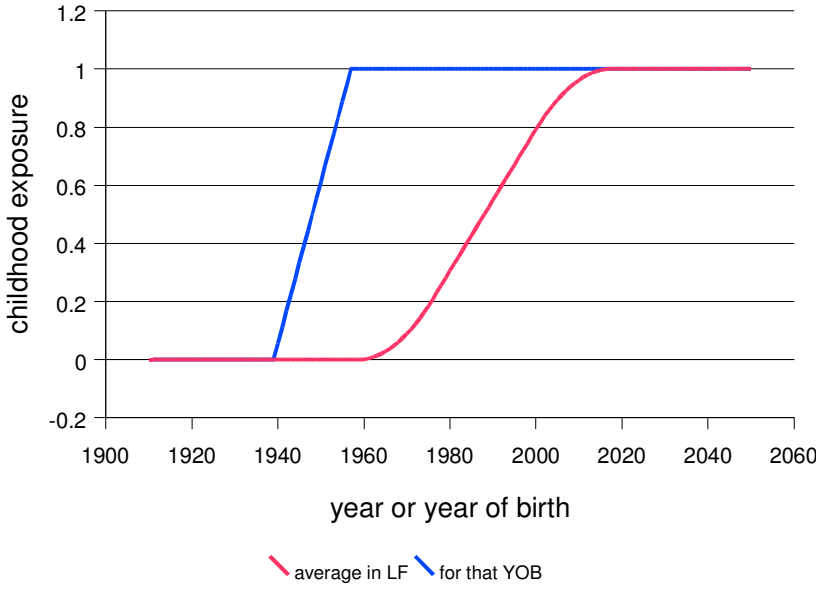

Figure 3: Malaria versus Childhood Exposure to Eradication, US States

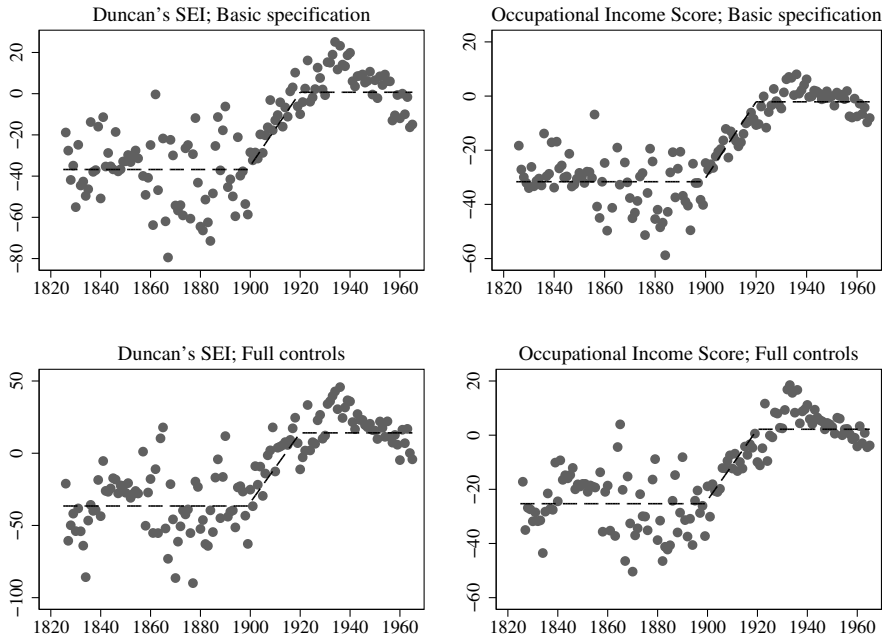

Notes: These graphics summarize regressions of income proxies on pre-eradication malaria-mortality rates. The y axis plots the estimated cohort-specific malaria coefficients. The x axis is the cohort's year of birth. The dashed lines measure the approximate childhood exposure to the malaria-eradication activities in the South. See Bleakley (2006b) for details on the data sources and regression specifications.

Figure 4: Adjustment to Mortality Declines: Output

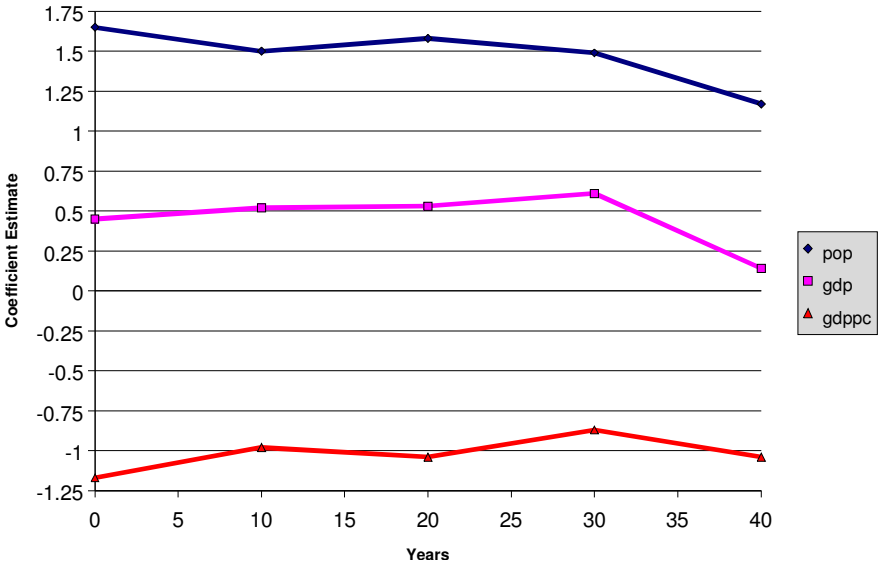

Notes: This graph reports 2SLS coefficients on life expectancy at various lags for the indicated outcomes. The estimates are drawn from various tables in Acemoglu and Johnson (2006).

Figure 5: Health and the Human Capital Decision