Volume Title: NBER Macroeconomics Annual 2008, Volume 23

Volume Author/Editor: Daron Acemoglu, Kenneth Rogoff and Michael Woodford, editors

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-00204-7

Volume URL: http://www.nber.org/books/acem08-1

Conference Date: April 4-5, 2008

Publication Date: April 2009

Chapter Title: Comment on "When Does Improving Health Raise GDP?"

Chapter Author: Hoyt Bleakley

Chapter URL: http://www.nber.org/chapters/c7279

Chapter pages in book: (p. 205 - 220)

Comment

Hoyt Bleakley, *University of Chicago and NBER*

I. Introduction

Under what circumstances does improving health raise productivity? How quickly does this happen? And when is this payoff large or just a fizzle? These are the questions addressed by Ashraf, Lester, and Weil. They are important questions without doubt. In the empirical literature to date, estimates of the effect of health on output range from the extremely positive to the substantially negative. Ashraf et al.'s approach—which has numerous merits to it, especially on incorporating population into the analysis—points us to a fairly compact middle ground on these estimates. This conclusion has a certain aesthetic appeal, but, in the end, I think that such precise predictions are difficult because of limitations of their simulation model and uncertainties about numerous parameters. I discuss my reasoning and provide a few empirical examples below.

The first theme of my comments is that health is multifaceted and that the type of health change matters for economic outcomes. Consider the following different improvements in health: better sanitation reduces infant mortality; eradication of parasitic disease reduces childhood morbidity (sickness short of death); low-cost HIV medicine shrinks adult mortality; an arthritis treatment curtails old-age morbidity; and improved treatments for pneumonia reduce old-age mortality. These are very different types of health improvements, and, intuitively, they will have different types of economic effects. An attempt to model all of health as a single index will run into trouble at some point. Most of the authors' results, however, are predicated on such a single-index assumption, linking health and human capital changes via a series of reduced-form relationships. I am concerned that some of the undergirding reduced-form equations might buckle if we do not have the 100% right model of health.

All that said, the Ashraf et al. framework might give the correct answer for the "average" disease, but most diseases are not average. This last detail matters. Because of this heterogeneity, there is probably a much larger range of possible outcomes from health improvements than the bounds that the authors simulate in their "best-case" and "worst-case" scenarios. Moreover, the most obvious policy levers to push on are not ones that improve general health, but instead are targeted at particular problems. Just as how Mister Wong and Professor Viner got to choose their lowest cost curve, so we could opt to intervene in a way skewed toward health improvements with higher payoffs. This means that the average returns from realized interventions would be higher than the return from treating the average disease.

In this vein, I consider the specific case of malaria. I argue that the authors' simulation understates the long-run effect of malaria on income by at least an order of magnitude. First, while the acute symptoms of malaria are measured by Ashraf et al., the disease is also characterized by long bouts of subclinical morbidity that are poorly measured. Second, morbidity during childhood might depress human capital formation, but the authors force this effect to work via years of schooling, which misses a lot and is itself a channel of second-order importance. Third, I review some of my recent estimates of the impact of childhood malaria exposure on adult income and find them to be considerably larger than the authors' simulations of the whole impact of malaria.

Finally, I discuss Ashraf et al.'s results on population. The expected response of population to health depends crucially on (*a*) whether the improvement is for mortality or morbidity and (*b*) how fertility responds to health. Their approach to fertility adjustment is seemingly comprehensive, but I provide below several examples of reasonably well-identified studies that give results outside the range of their simulations. I suggest an alternative framework but argue that further research is needed. Second, I discuss some limitations to the authors' Malthusian approach to population and income, especially in an urbanized and progressively integrating world. Third, I review some studies that estimate general equilibrium effects of health and find that the results of two out of three of them lie outside the range of their best/worst simulations.

II. Points of Agreement

But before I get on to the criticism, let me set the stage by reviewing some substantial points of agreement.

First, as Ashraf et al. point out, much of the effect of health on income will work through early-life exposure to disease, and this implies that there will be long lags before the effect on income is felt. An increase in child health could improve physical and cognitive development in childhood, but we have to wait for those cohorts to enter the labor force. This transition can take the better part of a century.

Second, time lost to adult disability in poor countries is nontrivial but small relative to the income gaps between rich and poor countries. The direct cost of disability—idleness from disease—is a small fraction of work hours. However, we know little about indirect effects of adult disability, such as ones that reduce the quality of labor input.

Third, when the health environment changes, so will the size and composition of the population. On the one hand, the effective supply of labor will increase: reducing mortality will increase the supply of warm bodies, whereas reducing morbidity will increase the amount those warm bodies can work. On the other hand, the returns to different types of skills will change, which in turn will alter people's human capital investment decisions. These resulting shifts in quantities supplied should all change relative prices. Accordingly, we need a general equilibrium model. Ashraf et al. use a standard neoclassical framework, combining a demographically driven change in effective labor with both a fixed factor and a slowly adjusting capital stock. This seems like the right way to start.

Fourth, the simple correlation between health and economic growth (or level of output) such as that present in cross-country comparisons is hard to interpret as the causal effect of health if micro or cohort-based studies are used as the guide. The authors support this view, although it bears mentioning that their model is neoclassical, in which health enters (directly or indirectly) as a level shift in the supply of effective labor. Growth is ruled out by assumption, except as part of a transition path to a new level following an improvement in health. What about health and the level of output? I agree with the authors that the micro/cohort evidence for causal effects of health on output falls short of the magnitude implied by a bivariate cross-country regression. How far short it falls is nevertheless something we disagree on.

Here endeth the agreeable part. In my judgment, the existing micro literature does point to effects of (some types of) health on output that are modestly sized. Note that I mean "modest" from the point of view of a macro/development economist: the income gap between Zambia and Sweden, for example, is very large, and the micro estimates of health effects explain only a small fraction of this difference. However,

from the point of view of a micro/labor economist, some of these estimates (including a few of my own) can seem quite large. From a policy perspective, some of these policy interventions look to me like low-hanging fruit.

III. Malaria and Beyond

In my judgment, the simulation by Ashraf et al. understates the long-run benefits of eradicating malaria by at least an order of magnitude. I discuss their simulation of the impact of malaria and how their main assumptions provide a limited picture of the disease's economic impact. Below, I review evidence that just one channel—childhood exposure—has an effect on adult income that is about a factor of 10 larger than the authors' estimate. Because they make use of a hypothetical malaria eradication in Zambia as a test bed for their simulation, I do so as well.[1]

One issue is that malarial fevers are only part of the morbidity associated with the disease. The benchmark used by the authors comes from the Global Burden of Disease project, which treats the prevalence of malaria as being the incidence of malarial fevers times the average fever duration (33% per year for adults and a few days, respectively, in Zambia, according to Ashraf et al.). The product of these two numbers is not very big. In words, fevers do not leave people so knocked out that they cannot work, except for a small fraction of the time. This datum is reminiscent of numerous studies of malaria from the 1960s that used time-use surveys to estimate the work time lost from experiencing a bout of malarial fever at a given moment. The conclusion of this literature was that the direct, contemporaneous effect of fevers on productivity was low (of order 1%). What is missing from this picture is that fevers are simply the most acute realization of sickness from malaria (the tip of the iceberg, as it were). Persistent infection from malaria (and almost all infection from malaria is persistent) causes considerable subclinical morbidity, including anemia.[2]

Of particular concern is the childhood morbidity associated with malaria, which has been shown to have particularly strong and long-lasting effects. Among children, malaria has been associated with stunting of physical and cognitive development, and it is not a leap to imagine that this reduces the return to human capital investment. Further, these damages would be hard to undo later on. But how to estimate the later-life impact?

One strategy to estimate such impacts is to use years of schooling, as the authors do, but this would yield an incomplete picture. As they

note, malaria's effect on education (estimated in Bleakley [2007b]) and Lucas [2007a]) is much larger than what we would predict on the basis of only its effect on longevity, in part because of its substantial child morbidity. But how much does childhood sickness affect income through education? Consider the individual's discounted lifetime income, y, at the optimal choice of schooling, $y^* \equiv y(e^*, h)$, where e is years of education and h is health (thought of here as less morbidity in childhood). This will respond to health via two channels, as seen by taking the full derivative of y^* with respect to h:

$$\frac{dy^*}{dh} = \frac{\partial y}{\partial e}\Big|_{e^*} \frac{de^*}{dh} + \frac{\partial y}{\partial h}\Big|_{e^*}.$$

The first term values the increase in years of schooling (de^*/dh) at the marginal return to schooling ($\partial y/\partial e|_{e^*}$). But notice that, by the envelope condition, the marginal return to schooling (for lifetime income) is zero at the optimal choice of schooling. So, changes in the quantity of education (i.e., time in school) are not of first-order importance and should essentially wash out to zero.

The second term measures the direct effect of health on labor productivity, evaluated at e^*. It is instructive to decompose this direct effect into two parts:

$$\frac{\partial y}{\partial h}\Big|_{e^*} = \frac{\partial y}{\partial h}\Big|_{e=0} + \int_0^{e^*} \frac{\partial^2 y}{\partial h \partial e}\, de.$$

The first part is the effect of health on income for those with no education ($e = 0$), and the second part is the changing returns to inframarginal schooling investments. These latter terms point to first-order benefits of health by raising the inframarginal return to education. In words, childhood health can raise the quality of education received (children can learn better, e.g.), and this can produce an increase in income of first-order importance.

In my own work (Bleakley 2007b), I attempt to directly estimate the effect of one's childhood exposure to malaria on income later in adulthood. I analyze malaria eradication campaigns in the United States (circa 1920)[3] and in Brazil, Colombia, and Mexico (circa 1955). Examining these episodes has a few useful features. On the one hand, these countries have (i) nonmalarial areas that can serve as a comparison group and (ii) census micro data available that cover the relevant sets of cohorts. On the other hand, these campaigns began because of advances in health technology that originated outside the affected regions, which reduces concerns about reverse causality. The basic finding of the study

is that cohorts born after eradication had higher income (and literacy) as adults than the preceding generation. This is true both in absolute terms and when measured relative to comparable cohorts in nonmalarious areas. Further, the timing of things supports my study's hypothesis as well: the observed changes coincide with childhood exposure to the campaigns rather than with preexisting trends. The magnitude of the malaria-induced income growth is nontrivial.

How do these estimates compare with Ashraf et al.'s simulated malaria eradication in Zambia? From my study, I chose two regions that were the most malarial in their respective countries: the Mississippi Delta states in the United States and the northern third of Brazil. I limit myself to these two examples in the interest of space. Reduced-form estimates for Mexico and Colombia are similar to those for Brazil, and magnitudes of income shifts are similar across all four countries when considered on a per-infection basis.

Some comparisons of relative disease rates in the three areas/times are found in panel A1 of table 1. First, we see that malaria accounts for approximately 10% of deaths in both Zambia today and the Mississippi Delta in 1890. (Such data for Brazil were not available for circa 1950.) In contrast, the overall death rate is not quite twice as high in Zambia as in the delta in 1890, so the malaria death rate per capita is quite a bit higher in Zambia. Next, I compare estimates of childhood malaria infection in these areas prior to malaria control efforts. (See my paper for details on how these are computed.) By these calculations, childhood malaria infection rates were about one-third in Mississippi Delta states and two-thirds in northern Brazil and Zambia. Note that this is a point-in-time probability referring to a stock of infection. Finally, the authors report the incidence of fevers among adults to be 33%, which is an annual probability referring to a flow of new fevers. Because (i) infection continues after an episode of fevers and (ii) children are more susceptible, this number is a lower bound for the static, childhood infection rate. Bottom line: malaria infection rates today in Zambia are higher than in the Mississippi Delta in 1890; instead Zambia's malaria rate is probably closer to northern Brazil's in 1950.

While the malaria burden is roughly similar across these three examples, Ashraf et al.'s simulated impact of malaria eradication is markedly different from the estimates using actual eradication campaigns. These results are found in panel A2 of table 1. Bear in mind that the estimated impacts from my study reflect one channel (childhood exposure to malaria), whereas the authors' simulated impacts purport to reflect all channels through which malaria depresses human capital. The first

Table 1
Impacts of Malaria and Hookworm Eradications

	A. Malaria		
	Mississipi Delta, 1890	Northern Brazil, 1950	Zambia, Present Day
A1. Malarial disease rates:			
Fraction of deaths due to malaria	10%	…	8.3%
Deaths per 1,000 population, per year	12	…	20
Precampaign malaria infection rate among children, point in time	33%	67%	67% (c. 1950)
Incidence of cases, adults, per year	…	…	33%
A2. Long-run impacts:			
Increase in human capital (partial equilibrium)[a]	10%–15%	20%–30%	(?)2.75%
Increase in population	c. 2.5%	c. 5%	5%
Increase in income/worker, adjusted for land dilution (general equilibrium)	8%–13%	16%–25%	1.15%

	B. Hookworm	
	Southern United States, 1910	Kenya, Present Day
B1. Estimates:		
Childhood infection rate	40%	90%
Drop in adult income from persistent childhood exposure (partial equilibrium)	43%	…
B2. Long-run impacts:		
Average gain in income/worker from eradication (partial equilibrium)	17%	39%
Increase in income/worker adjusted for land dilution (general equilibrium)	15%	35%

Note: Panel A considers malaria eradication campaigns in the United States and Brazil and a hypothetical campaign in Zambia in the present day. For panel A1, data for the United States and Brazil are drawn from Bleakley (2007b), as is the Zambian estimate for precampaign childhood infection. The remaining Zambian numbers in panel A1 are taken from Ashraf et al.'s estimates. The estimated change in the Zambian population in panel A2 comes from their fig. 21. The corresponding Brazilian number is assumed to be similar, whereas the U.S. number is set as half the Zambian value, which takes into account approximate differences in malaria rates among these three countries. The malaria eradication–induced increase in "human capital" for the United States and Brazilian regions is based on estimates using comparisons of income across cohorts in Bleakley (2007b). The corresponding (but not strictly comparable) Zambian number is computed by Ashraf et al., who consider contributions to income from lower disability and more years in school only. The second row in panel A2 presents simulations of the long-run change (relative to the baseline) in population following malaria eradication. The number for Zambia is computed by Ashraf et al. using a cause-deleted life table; the numbers for the United States and Brazil are scaled versions of the Zambian number using estimates of precampaign malaria infection among children. The final row in panel A2 adjusts the previous numbers for the effect of land dilution because of both the increase in population and effective labor supplied. These adjustments follow the Ashraf et al. methodology, which assumes that different skill types are perfect substitutes and a land share of 0.1. Panel B considers the antihookworm campaign in the southern United States and a hypothetical campaign in Busia, Kenya. Estimates and extrapolations are drawn from Bleakley (2007a), except for the final row in panel B2, which is computed as in panel A2 above. Blank cells indicate a lack of data.
[a]Figures for the Mississippi Delta and northern Brazil come from Bleakley (2007b); the figure for Zambia comes from Ashraf et al.

row of panel A2 displays the estimated partial equilibrium impact on human capital. Is it plausible that Zambia would have a response to malaria eradication that is an order of magnitude below what was estimated for northern Brazil at midcentury, a region with a comparable malaria burden? Further, is it not even more surprising that Zambia's simulated response is only a quarter of what I estimate for the circa 1920 Mississippi Delta, an area with an undoubtedly lower malaria intensity? The estimated changes in income for the United States and Brazil, as well as results for Colombia and Mexico, were roughly in proportion to their relative malaria burdens. So why would we expect Zambia to be so far out of line?

At this point, there seem to be two possibilities: (i) my estimates are bunk, albeit in a way that is common across all four countries I study; or (ii) there is something about malaria that makes it incompatible with the authors' calculation. Needless to say, I do not favor option i, although I invite the reader to review the above-referenced paper and decide for himself or herself. However, option ii seems probable on the basis of the presentation above about the incomplete measurement of malaria morbidity and the inappropriateness of quantity of schooling as a unitary channel for human capital effects.

In an important contribution, Ashraf et al. show how the increase in effective labor from malaria eradication will depress productivity in a land-dependent economy, so I correct for this in the remainder of panel A2. Their simulated population change for Zambia is scaled linearly for the other two regions, on the basis of the estimated infection rates prior to malaria control. The simulated population change is combined with the estimated (or simulated) change in human capital to compute the change in effective labor. (These calculations, like those in Ashraf et al.'s paper, treat different skill types as perfect substitutes, which might be a deficiency. See below.) As before, these results show a stark contrast between the authors' simulation and my estimated responses to malaria eradication.

I am not inclined to believe that the result for malaria and morbidity is an aberration, because this pattern is also seen in my earlier work on hookworm disease (Bleakley 2007a). Hookworm, a parasite that infects humans, is like malaria in that it has a large burden of childhood morbidity (including anemia), but with considerably less mortality. The children of the southern United States had hookworm infection rates of 30%–40%, and areas of the sandy coastal plain had close to 100% infection rates among kids. In the 1910s, John D. Rockefeller, being the Bill Gates of his day, donated a considerable sum of money to a campaign

to "deworm" the South. Before then, the disease was mostly unheard of and was never treated.

Results of the antihookworm campaign were remarkable. The effect of deworming children was illustrated in testimonials from across the region, such as this one from Varnado, Louisiana: "In short, we have here in our school-rooms today about 120 bright, rosy-faced children, whereas had you not been sent here to treat them we would have had that many pale-faced, stupid children" (Rockefeller Sanitary Commission 1912, 58).

Moving beyond anecdotes, I used census data to systematically examine the impact of hookworm eradication on children. Similar to the malaria study cited above, areas that were hookworm free served as a comparison group. Children growing up in areas that benefited from the antihookworm campaign saw large increases in literacy and income, relative to earlier cohorts.

According to my estimates, childhood exposure to hookworm depressed adult income by a lot in the South, and this result suggests that large gains from mass deworming are possible in Africa. Pertinent numbers are seen in panel B of table 1. As mentioned above, children in the South had a 40% hookworm infection rate circa 1910, which compares with a 90% childhood infection rate from hookworm in the Busia region of Kenya.[4] My instrumental variables estimates of the effect of hookworm on income can be interpreted as follows: If I take your point-in-time probability of hookworm infection in childhood from zero to one, it reduces young adult income by 43%, in natural log terms. (Note that this refers to persistent infection, not to having ever been infected.) Eradication would therefore imply a long-run human capital gain of 17% in the American South. Extrapolating to Busia, we would expect a 39% long-run increase in human capital from eradicating hookworm. Because hookworm disease has little associated mortality, eradication would imply negligible direct changes in population, although effective labor would rise with the increase in human capital. Using the authors' general equilibrium adjustment, I compute the simulated changes in income per worker, which are a bit lower than the human capital estimates. (Results are found in panel B2 of table 1.) In any event, all these numbers point to large income gains from hookworm eradication.

Estimates from these eradication campaigns also illustrate the problem with using the quantity of education as the chief channel of childhood morbidity effects on income. In none of the cases studied do years of schooling account for more than a quarter of the effect of income.

Nor should we expect time in school to be the central channel, as per the argument above. Indeed, in one case (malaria eradication in Mexico), adult income rises with less childhood exposure to malaria, but years of school actually decrease! It is evident that treating the quantity of education as a central outcome is misleading when analyzing certain changes in health.

Are hookworm and malaria unusual diseases? No, in the sense that over a billion people worldwide are infected with these and related parasites. Yes, in the sense that these and related parasites cause a relatively high burden of morbidity, especially among children. These parasites tend to persist in the system for a long time, since they have evolved to fly under our immune system's radar. These features make for a sharp contrast with diseases that have high case fatality rates and/or short-lived bouts of morbidity. As a point of contrast, consider chicken pox. Vaccination may just eradicate the disease starting with my daughters' generation, but I myself did not have this prophylaxis growing up. I spent a week of my childhood at home with the disease and watched baseball on television. The long-term consequences for me were probably not that great: a week of lost school, which I partially made up for by working harder in the subsequent weeks. Smallpox is another example: short bouts of morbidity and a high case fatality rate; but generally the residual disfigurement among the survivors does not cause a high disability burden. Whatever effects eradicating smallpox had, they probably worked through changes in mortality rates, and the consequences were therefore very different from the effects of malaria eradication. It should be clear by now that I do not argue (nor do I believe) that the eradication of any conceivable disease would necessarily have such large effects as those seen for hookworm and malaria. But unpackaging health seems key.

The evidence from historical campaigns against malaria and hookworm brings us back to the Abuja declaration, which Ashraf et al. used to motivate their study. The Abuja document states that income in Africa would be 37% higher today but for malaria. Results from the Americas suggest that malaria eradication would increase Zambian income per capita by 10%–30% in the long run, and Zambia is not the most malarial country in Africa. Moreover, this estimate is just for the gains realized via the mechanism of reduced childhood exposure to the disease. Presumably there are other benefits as well. As the authors mention, a significant cost of malaria might be that investment in physical capital and land improvement are suppressed by the threat of the disease. For example, the control of malaria in the Panamanian isthmus

was crucial for the completion of the canal by the United States, and the earlier failure of the same French effort was partly due to a lack of scientific knowledge about malaria transmission at the time. Meanwhile, back in Zambia, Utzinger et al. (2001) argue that the localized control of malaria transmission was a key factor in the development of copper mining in that country. Finally, Bleakley and Hong (2008) relate agricultural productivity in the southern United States to the decline of malaria and find large effects. These effects, like those linked to childhood exposure above, would nevertheless take some time to kick in.

IV.   Whither Population?

A central contribution of Ashraf et al.'s paper is quantifying how much health improvements will increase population. Increases in population in large measure stem from people not dying as quickly as before. In contrast, reductions in morbidity (and increases in human capital) do not increase population, but do raise the supply of effective labor to a lesser extent. Again, the type of health shock matters for the result.

Nevertheless, for a given change in health, the long-run change in population depends crucially on the response of fertility, as the authors ably demonstrate. Unfortunately, our understanding of fertility's adjustment to large changes in health (among other things) leaves something to be desired. As they point out, there has been enormous heterogeneity in the demographic transition (or lack thereof) both historically in the developed world and in recent generations in less developed countries. This sort of instability in a reduced-form relationship highlights the complicated nature of the comovement between health and fertility.

Perhaps unpackaging our notion of health a bit can aid in understanding the fertility response. Consider three recent studies of fertility, placed in the framework of the quantity/quality (q^2) trade-off. Acemoglu and Johnson (2007) examine the large declines in mortality (especially among the very young) that took place circa 1950 and find large, positive responses in birth rates. Their shock is like a decline in the price of quantity ($p_n \downarrow$) in the q^2 model. Accordingly, quantity rises in their estimates. (Note that slow, monotonic adjustment to a new fertility level would yield essentially the same birth rate after the shock as before.) In contrast, Fabian Lange and I (Bleakley and Lange 2009) examine the fertility response to hookworm eradication in the southern United States. Getting rid of hookworm in effect reduced the price of quality because it was now easier to rear an educated child. We found that fertility goes down with

the reduction in hookworm infection, which is again consistent with the q^2 model.

What about when both prices change? Consider the work of Adrienne Lucas (2007a, 2007b), who estimates fertility rates before, during, and after the malaria eradication campaign (circa 1950) in Sri Lanka. Her results vary by generation. Women in their fertile years around the time of the campaign had higher fertility following the decline in malaria, but women in childhood at that time went on to have fertility rates comparable to those in the nonmalarious parts of Sri Lanka.[5] Both generations had a lower p_n than previous cohorts, but the second generation, which had escaped exposure to malaria in childhood, also had higher human capital.

So, we have three improvements in health and three distinct responses. I believe that all three lie outside the range of what the authors' fertility model would have predicted. In contrast, this diversity of effects seems to make sense in a simple q^2 model, although other interpretations are possible. What is clear, however, is that these results are inconsistent with a univariate model of health. Further research is needed to better understand the effects of health on fertility.

V.   Why Just Malthus?

The improvements in health contemplated in this literature (including the paper by Ashraf et al.) are huge, and it seems more than likely that there would be general equilibrium effects. An important contribution of this paper is to consider how much the population increase dilutes land and capital. But why restrict ourselves to Malthusian effects when presumably other prices will change as well?

For the moment, consider first the general equilibrium effects of just population increasing. If the working-age population rises by x%, does the average plot of land get x% more farmers? Recent historical experience suggests otherwise. Latin America, for example, went from 40% urban at mid-twentieth century to almost 80% today. Further, the majority of the world's population lives in urban areas as of a few years ago. This situation seems less like an economy that is land dependent and decreasing returns to scale in reproducible factors. Instead, it seems like one in which agglomeration economies might kick in. (See Jena et al. [2007] for more analysis on the point.)

How does it modify the results if all the extra population goes to the city? In the short run, we should obtain similar capital dilution and compositional change in the population. In the medium run, fertility will

adjust, although the model of fertility is presumably a different animal when the economy has both rural and urban sectors. Finally, in the long run, land is much less of a constraint than in the Ashraf et al. simulations. (Even in a highly urbanized economy, congestion effects will remain in one form or another, albeit ones that kick in at a much higher population density.)

Beyond just population, health improvements will change the mix of skills in the labor force. Ashraf et al. use a simplifying assumption that different types of skills—health, schooling, and experience—are perfect substitutes. While we do not know much about these parameters for poor countries, estimates in the labor literature for the United States contradict this assumption.

Relaxing the assumption of perfect substitution among skill types will tend to magnify the range of uncertainty in their simulations, and again this will depend on which element of the health vector is being changed. Consider a few cases. First, suppose that the shock to health brings more human capital to an economy in which skilled labor is very scarce. In this case, the response of output would be more positive than in the Ashraf et al. baseline. Suppose instead that the change in health did little but decrease mortality among infants in poor families. This would likely increase the supply of low-skilled labor in an economy that is already skill scarce and depress average income even more than the authors' baseline.

Finally, the assumption of a closed economy seems increasingly unrealistic for the twenty-first century. How would openness change the results? In their sensitivity analysis, the authors allow for capital to flow in from the outside world. Not surprisingly, this reduces capital dilution, and therefore output per capita falls less in the short run when population rises. Next, openness to international trade would surely modify the authors' results. If the economy is reasonably diversified, we would expect the price responses to health changes to be attenuated (and, in the extreme, we would expect factor price equalization). Finally, migration can act as a "safety valve" when an area becomes overpopulated. For example, Hanson and McIntosh (2007) find that relatively large birth cohorts in Mexico had higher propensities to migrate to the United States.

The empirical literature analyzing general equilibrium effects of health shocks provides a range of estimates that is even larger than that of Ashraf et al.'s sensitivity analysis. Perhaps most consistent with their results is a classic study by Schultz (1964). He considers the impact of the 1918 influenza pandemic on farm output in India. The flu killed substantial numbers in India but left capital and land intact, and it also had little scarring effect on adults who survived. Schultz therefore

interprets the flu mortality as a decline in population only. Comparing a few years before to a few years after the epidemic (i.e., the short run), he measures declines in output that were greater in areas that suffered more flu deaths. The magnitude of this relationship was consistent with a labor share of around 0.5. This result is in line with those of Ashraf et al., at least for the short run.

In contrast, two recent studies find markedly different responses of aggregate output to health shocks. In cross-national data, Acemoglu and Johnson (2007) find that, following a decline in mortality, there was an increase in population but no statistically significant rise in GDP. These results are qualitatively similar to Ashraf et al.'s "worst-case" scenario in which fertility adjusts very slowly and the land share is high. Quantitatively, however, Acemoglu and Johnson's estimated population change is considerably higher than any that Ashraf et al.'s model can generate, and, moreover, the decline in GDP per capita is much larger than they would predict, even accounting for the higher-than-simulated increase in population. In my own work (Bleakley 2007c), I estimate the aggregate (i.e., state-level) responses to the above-mentioned campaigns against hookworm and malaria in the United States. Areas that stood to benefit from eradication saw slow increases in output that tracked the entrance into the workforce of cohorts born late enough to have escaped childhood exposure to these diseases. Such long lags are consistent with Ashraf et al.'s simulations in which the health effects on productivity operate at the cohort level (via early-life health) rather than contemporaneously with the change in health. My estimated magnitudes are, if anything, larger than what I obtain in the earlier cohort-level estimates (which are already larger than what the authors simulate, as I discuss above). At the risk of sounding like a broken record, the contrasting results of these two studies seem linked to the different sorts of changes in health considered, with the former tilted toward mortality (especially infant) and the latter toward childhood morbidity.

VI.   Concluding Thoughts

What is the motivation of economists in studying interrelationships between health and development? Is the goal to understand whether health is the single factor that explains everything, leading to a "magic bullet" intervention that fixes everything? Are we instead looking for doable interventions that pass the cost/benefit test? If it is the former, health—however defined—is certain to fail. If it is the latter, then we

should turn our attention to the numerous examples of successful campaigns, some of which are based on century-old technology.

One important issue that remains is whether it is feasible to undertake such interventions where they are most needed. Above, I discussed Rockefeller's campaign against hookworm in the American South, but an interesting episode for comparison comes from Puerto Rico. Around the same time, a commission from the U.S. Army sponsored an antihookworm campaign throughout that Caribbean island. Large gains against hookworm were realized immediately after the campaign. Unfortunately, the colonial government provided very little follow-up support, and these gains had mostly disappeared a decade later. Improving health is a long, hard slog, and potential interventionists should be prepared for as much.

Endnotes

I thank Jane Fortson and Miriam Wasserman for helpful discussions.

1. The other disease-specific case the authors consider is tuberculosis. This is a disease whose effects occur mostly in adulthood and whose morbidity is reasonably well measured. As a consequence, it might be well approximated by the sum of the various channels articulated by the authors.

2. Using a randomized trial of iron supplements in Indonesia, Thomas et al. (2003) show that anemia depresses contemporaneous adult productivity.

3. So-called tropical diseases, such as malaria and hookworm, were present in subtropical regions a century ago and were in many cases the focus of "big push" campaigns for eradication.

4. Data on hookworm infection for Zambia were not readily available. Busia and Zambia are ecologically similar to some degree (similar altitudes and temperatures, both proximate to great lakes). However, Zambia is a bit more rich and urbanized and has more seasonal rainfall, all of which would suggest a lower hookworm infection rate in Zambia.

5. These three studies of fertility all use never-infected or minimally infected areas as comparison groups, and statements about the effects of health changes are estimated relative to the control areas. Such areas did not benefit directly from the health improvements, but would have been exposed to many of the other shocks that prevailed in the episodes studied.

References

Acemoglu, Daron, and Simon Johnson. 2007. "Disease and Development: The Effect of Life Expectancy on Economic Growth." *Journal of Political Economy* 115 (December): 925–85.

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

———. 2007b. "Malaria in the Americas: A Retrospective Analysis of Childhood Exposure." Working paper, Graduate School of Business, University of Chicago.

———. 2007c. "Spillovers and Aggregate Effects of Health Capital: Evidence from Campaigns against Parasitic Disease in the Americas." Manuscript, University of Chicago.

Bleakley, Hoyt, and Sok Chul Hong. 2008. "The Impact of Weather on U.S. Farm Productivity: Historical Patterns and Relation to the Changing Disease Environment." Manuscript, Center for Population Economics, University of Chicago.

Bleakley, Hoyt, and Fabian Lange. 2009. "Chronic Disease Burden and the Interaction of Education, Fertility and Growth." *Review of Economics and Statistics*, forthcoming.

Hanson, Gordon H., and Craig McIntosh. 2007. "The Great Mexican Emigration." Working Paper no. 13675 (December), NBER, Cambridge, MA.

Jena, Anup, Casey Mulligan, Tomas Philipson, and Eric Sun. 2007. "The Value of Life in General Equilibrium." Manuscript, University of Chicago.

Lucas, Adrienne M. 2007a. "The Impact of Disease Eradication on Fertility and Education." Manuscript, Wellesley College.

———. 2007b. "The Impact of Malaria Eradication on Zambian Fertility: A Multigenerational Projection." Manuscript, Wellesley College.

Rockefeller Sanitary Commission. 1912. *Annual Report*. New York: Rockefeller Foundation.

Schultz, Theodore W. 1964. *Transforming Traditional Agriculture*. Studies in Comparative Economics, no. 3. New Haven, CT: Yale University Press.

Thomas, Duncan, et al. 2003. "Iron deficiency and the Well-Being of Older Adults: Early Results from a Randomized Nutrition Intervention." Manuscript, University of California, Los Angeles.

Utzinger, Jürg, Yesim Tozan, Fadi Doumani, and Burton H. Singer. 2001. "The Economic Payoffs of Integrated Malaria Control in the Zambian Copperbelt between 1930 and 1950." *Tropical Medicine and International Health* 7 (August): 657–77.